



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **Cisco MDS 9000 Family Fabric Manager Configuration Guide, Release 2.x**

Cisco MDS SAN-OS Release 2.0(1b) through Release 2.1(1a)  
March 2005

### **Corporate Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 526-4100

Text Part Number: OL-6965-02



THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Access Registrar, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StrataView Plus, TeleRouter, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0502R)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **New and Changed Information   xxv**

### **Preface   xxxi**

Audience	xxxi
Organization	xxxi
Document Conventions	xxxiv
Related Documentation	xxxiv
Obtaining Documentation	xxxv
Cisco.com	xxxv
Documentation DVD	xxxv
Ordering Documentation	xxxvi
Documentation Feedback	xxxvi
Cisco Product Security Overview	xxxvi
Reporting Security Problems in Cisco Products	xxxvii
Obtaining Technical Assistance	xxxvii
Cisco Technical Support Website	xxxvii
Submitting a Service Request	xxxviii
Definitions of Service Request Severity	xxxviii
Obtaining Additional Publications and Information	xxxix

---

## **PART 1**

---

## **Fabric Manager Applications**

---

### **CHAPTER 1**

### **Installation and Configuration   1-1**

About Cisco Fabric Manager	1-1
Fabric Manager Server	1-2
Fabric Manager Client	1-2
Fabric Manager Server Proxy Services	1-2
Device Manager	1-2
Performance Manager	1-3
Fabric Manager Web Services	1-3
Cisco MDS 9000 Switch Management	1-3
Storage Management Solutions Architecture	1-4
In-Band Management and Out-of-Band Management	1-5
mgmt0	1-5

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**IPFC 1-5**

Installing the Management Software	1-5
Before You Install	1-6
Installation Procedure	1-7
Upgrading the Management Software	1-8
Downgrading the Management Software	1-9
Downgrading to Cisco MDS SAN-OS Release 2.x or Later	1-9
Downgrading to Cisco MDS SAN-OS Release 1.3(x) or Earlier	1-9
Launching the Management Software	1-9
Integrating Cisco Fabric Manager with Other Management Tools	1-11
Running Fabric Manager Behind a Firewall	1-11
Uninstalling the Management Software	1-12

---

**CHAPTER 2**

**Fabric Manager Server 2-1**

Fabric Manager Server Overview	2-1
Fabric Manager Server Features	2-2
Installing and Configuring Fabric Manager Server	2-2
Installing Fabric Manager Server	2-3
Unlicensed Versus Licensed Fabric Manager Server	2-3
Setting the Seed Switch	2-4
Configuring Flows and Collections with Performance Manager	2-4
Using the Performance Manager Configuration Wizard	2-4
Installing Fabric Manager Web Services	2-6
Verifying Performance Manager Collections	2-6
Fabric Manager Server Fabric Monitoring and Removal	2-7
Designating a Fabric for Continuous Monitoring	2-7
Removing a Fabric from Monitoring	2-8
Fabric Manager Server Properties File	2-8

---

**CHAPTER 3**

**Fabric Manager Client 3-1**

Fabric Manager Client Overview	3-1
Fabric Manager Advanced Mode	3-2
Launching Fabric Manager Client	3-2
Using Fabric Manager Client	3-3
Multiple Fabric Display	3-4
Contents Panes	3-4
Fabric Pane	3-5
Main Menu	3-7

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Toolbar	3-7
Information Pane	3-9
Logical Domains Pane	3-10
Physical Attributes Pane	3-10
Status Bar	3-10
Context Menus	3-11
Filtering	3-11
Detachable Tables	3-12
Setting Fabric Manager Preferences	3-12
Network Fabric Discovery	3-14
Modifying Device Grouping	3-14
Control of Administrator Access with Users and Roles	3-15
Fabric Manager Wizards	3-15
Fabric Manager Troubleshooting Tools	3-16

## CHAPTER 4

<b>Device Manager</b>	<b>4-1</b>
Device Manager Overview	4-1
Device Manager Features	4-1
Launching Device Manager	4-2
Using Device Manager	4-3
Menu Bar	4-4
Toolbar Icons	4-4
Dialog Boxes	4-5
Tabs	4-5
Legend	4-6
Supervisor and Switching Modules	4-7
Context Menus	4-7
Setting Device Manager Preferences	4-8

## CHAPTER 5

<b>Fabric Manager Web Services</b>	<b>5-1</b>
Fabric Manager Web Services Overview	5-1
Filter Tree	5-2
Events	5-3
Performance	5-4
Inventory	5-4
Admin	5-4
Installing Fabric Manager Web Services	5-5
Using Fabric Manager Web Services with SSL	5-6

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Launching and Using Fabric Manager Web Services 5-7
- Monitoring Fabrics from Fabric Manager Web Services 5-9
- Setting Up a Guest User 5-9

## CHAPTER 6

- Performance Manager 6-1**
  - Performance Manager Architecture 6-1
  - Data Interpolation 6-2
  - Data Collection 6-2
  - Using Performance Thresholds 6-2
  - Quick Data Collector and Flow Setup Wizards 6-3**

## CHAPTER 7

- Authentication in Fabric Manager 7-1**
  - Fabric Manager Authentication Overview 7-1
  - Best Practices for Discovering a Fabric 7-3
    - Setting up Discovery for a Fabric 7-3
  - Performance Manager Authentication 7-3
  - Fabric Manager Web Services Authentication 7-4

## CHAPTER 8

- Cisco Traffic Analyzer 8-1**
  - Using Cisco Traffic Analyzer with Performance Manager 8-1
    - Understanding SPAN 8-2
    - Understanding the PAA-2 8-3
    - Understanding Cisco Traffic Analyzer 8-3
  - Installing Cisco Traffic Analyzer 8-4
  - Configuring Performance Manager for Use with Cisco Traffic Analyzer 8-5

## PART 2

### **Cisco MDS SAN-OS Installation and Configuration Files**

## CHAPTER 9

- Obtaining and Installing Licenses 9-1**
  - Licensing Terminology 9-1
  - Licensing Model 9-2
    - Licensing High Availability 9-4
  - Options to Install a License 9-5
    - Obtaining a Factory-Installed License 9-5
    - Performing a Manual Installation 9-5
    - Obtaining the License Key File 9-6
    - Installing the License Key File 9-7
  - Installing Licenses Using Fabric Manager License Wizard 9-7

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Viewing License Information in Fabric Manager	9-8
Installing or Updating Licenses Using Device Manager	9-9
Viewing License Information in Device Manager	9-10
Uninstalling Licenses	9-10
Updating Licenses	9-11
License Expiry Alerts	9-11
Moving Licenses Between Switches	9-12
Fabric Manager Server Licensing	9-12

---

## CHAPTER 10

### Software Images 10-1

About Software Images	10-1
Dependent Factors	10-1
Essential Upgrade Prerequisites	10-2
Software Upgrade Methods	10-3
Determining Compatibility	10-3
Recognizing Failure Cases	10-4
Using the Software Install Wizard	10-4
Upgrading from Cisco MDS SAN-OS 1.3(4a) to 2.0(1b)	10-6
File System Manipulation	10-7
Listing the Files in a Directory	10-8
Creating a Directory	10-8
Deleting an Existing File or Directory	10-9
Copying Files	10-9
Performing Other File Manipulation Tasks	10-10

---

## CHAPTER 11

### Configuration Files 11-1

Working with Configuration Files	11-1
Saving the Configuration File	11-1
Copying the Configuration File	11-2

---

## PART 3

---

### Switch Configuration

---

## CHAPTER 12

### Cisco Fabric Services 12-1

About CFS	12-1
Cisco MDS SAN-OS Features Using CFS	12-1
CFS Features	12-2
Enabling CFS for a Feature	12-3

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Locking the Fabric	12-4
Committing Changes	12-4
Clearing a Locked Session	12-6
Disabling or Enabling CFS Distribution on a Switch	12-6
CFS Merge Support	12-7
A CFS Example Using Fabric Manager	12-7
A CFS Example Using Device Manager	12-9

## CHAPTER 13

### **VSAN Configuration** 13-1

About VSANs	13-1
Default and Isolated VSANs	13-1
Default VSANs	13-2
Isolated VSANs	13-2
Configuring a VSAN	13-2
Deleting VSANs	13-3

## CHAPTER 14

### **Dynamic VSAN Configuration** 14-1

About DPVM	14-1
DPVM Requirements	14-2
DPVM Databases	14-2
DPVM Database Distribution	14-2
Config Database Activation	14-3
Copying the DPVM Database	14-3
Autolearn Entries	14-3
Using the DPVM Setup Wizard	14-4
Modifying the DPVM Database	14-4
Using the DPVM tables	14-5

## CHAPTER 15

### **Zone Configuration** 15-1

Zoning Features	15-1
Zone Implementation	15-2
Zone Configuration	15-2
Using the Zone Configuration Tool	15-3
Edit Full Zone Database Overview	15-4
Zone Database Information	15-4
Configuring a Zone	15-5
Viewing Zone Statistics	15-5
Adding Zone Members	15-5



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Displaying Zone Membership Information	15-6
Alias Configuration	15-6
Creating Zones with Aliases	15-6
Viewing Aliases	15-7
Zone Set Creation	15-8
Active and Full Zone Set Considerations	15-8
Creating Zone Sets	15-11
Adding Zones to a Zone Set	15-11
Activating Zone Sets	15-11
Deactivating Zone Sets	15-12
Creating Additional Zones and Zone Sets	15-12
Cloning Zones and Zone Sets	15-12
Deleting Zones, Zone Sets, and Aliases	15-13
Zone Enforcement	15-13
The Default Zone	15-14
Configuring the Default Zone Policy	15-14
Performing Zone Merge Analysis	15-15
Recovering from Link Isolation	15-15
Importing Zone Sets	15-16
Exporting Active Zone Sets	15-17
Full Zone Set Propagation	15-17
One-Time Distribution	15-17
Copying a Full Zone Database	15-18
Migrating a Non-MDS Database	15-18
Zone-Based Traffic Priority	15-18
Configuring Zone QoS and Broadcast Attributes	15-18
About LUN Zoning	15-19
Configuring a LUN-Based Zone	15-20
Assigning LUNs to Storage Subsystems	15-21
About Read-Only Zones	15-21
Guidelines to Configure Read-Only Zones	15-21
Configuring Read-Only Zones	15-22
Backing Up and Restoring Zones	15-22

## CHAPTER 16

## Inter-VSAN Routing Configuration 16-1

Inter-VSAN Routing	16-1
Understanding IVR	16-1
IVR Terminology	16-2
Fibre Channel Header Modifications	16-3

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

IVR NAT	16-3
IVR VSAN Topology	16-4
Autonomous Fabric ID	16-4
Service Groups	16-4
Using IVR NAT and Auto Topology	16-5
Transit VSAN Guidelines	16-5
Border Switch Guidelines	16-5
Service Group Guidelines	16-6
Using IVR Without IVR NAT or Auto Topology	16-6
Domain ID Guidelines	16-6
Transit VSAN Guidelines	16-7
Border Switch Guidelines	16-7
Using the IVR Zone Wizard	16-7
Modifying IVR	16-8
Modifying IVR NAT and IVR Auto Topology	16-9
Configuring Service Group	16-9
Configuring AFIDs	16-9
Enabling IVR Without NAT	16-10
Manually Creating the IVR Topology	16-11
Activating an IVR Topology	16-12
Clearing the IVR Topology	16-12
Adding IVR Virtual Domains	16-12
IVR Zones and IVR Zone Sets	16-13
IVR Zones Versus Zones	16-13
Automatic IVR Zone Creation	16-14
Configuring IVR Zones and Zone Sets	16-14
Creating Additional IVR Zones and Zone Sets	16-15
Activating IVR Zone Sets	16-16
Deactivating IVR Zone Sets	16-16
Recovering an IVR Full Zone Database	16-16
Recovering an IVR Full Topology	16-17
Adding Members to IVR Zones	16-17
IVR Interoperability	16-17

## CHAPTER 17

### **PortChannel Configuration** 17-1

PortChannel Functionality	17-1
Using the PortChannel Wizard	17-2
Modifying PortChannels	17-5

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## CHAPTER 18

### Interface Configuration 18-1

#### Fibre Channel Interfaces 18-1

##### About Interface Modes 18-1

E Port 18-2

F Port 18-2

FL Port 18-3

TL Port 18-3

TE Port 18-3

SD Port 18-3

ST Port 18-4

Fx Port 18-4

B Port 18-4

Auto Mode 18-4

Configuring Trunking Mode 18-4

##### About Interface States 18-5

Administrative States 18-5

Operational States 18-5

Reason Codes 18-5

32-Port Configuration Guidelines 18-5

Configuring Fibre Channel Interfaces 18-6

Configuring Gigabit Ethernet Interfaces 18-7

Enabling or Disabling Interfaces 18-7

Managing Interface Attributes for Ports 18-7

Buffer-to-Buffer Credits 18-8

Performance Buffers 18-8

Configuring Buffer-to-Buffer Credits and Performance Buffers 18-8

Identification of SFP Types 18-9

Configuring the Management Interface 18-9

Configuring Persistent FC IDs 18-9

IPFC Interface Configuration 18-9

## CHAPTER 19

### FCIP Configuration 19-1

About Gigabit Ethernet Interfaces 19-1

Configuring a Basic Gigabit Ethernet Interface 19-2

FCIP Configuration 19-2

FCIP and VE Ports 19-2

FCIP Links 19-3

FCIP Write Acceleration 19-4

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

FCIP Compression	19-5
Using the FCIP Wizard	19-5
Modifying FCIP Links	19-8
About FCIP Profiles	19-8
FCIP Interfaces	19-9
Modifying FCIP Profiles and FCIP Links	19-9
Verifying Interfaces and Extended Link Protocol	19-10
Checking Trunk Status	19-10
Modifying FCIP Write Acceleration or FCIP Compression	19-11
FCIP Tape Acceleration	19-11
Enabling FCIP Tape Acceleration	19-13
Configuring Advanced FCIP Interfaces	19-13
Configuring Peers	19-13
Peer IP Address	19-13
Special Frames	19-14
Using B Port Interoperability Mode	19-15
Configuring B Ports	19-17
Configuring E Ports	19-18
FCIP High Availability	19-18
Fibre Channel PortChannels	19-19
FSPF	19-19
VRRP	19-20
Ethernet PortChannels	19-20
Ethernet PortChannels and Fibre Channel PortChannels	19-21

## CHAPTER 20

<b>iSCSI Configuration</b>	20-1
Configuring iSCSI	20-1
About iSCSI	20-1
Routing iSCSI Requests and Responses	20-4
Enabling iSCSI	20-5
Using the iSCSI Wizard	20-5
Presenting Fibre Channel Targets as iSCSI Targets	20-7
Dynamically Importing Fibre Channel Targets	20-8
Creating a Static iSCSI Virtual Target	20-9
High Availability Static Target Importing	20-10
Configuring the Trespass Feature	20-11
Presenting iSCSI Hosts as Virtual Fibre Channel Hosts	20-11
Dynamic Mapping	20-12
Static Mapping	20-12

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Assigning VSAN Membership to iSCSI Hosts	20-13
Creating a Statically Mapped iSCSI Initiator	20-13
iSCSI Proxy Initiators	20-14
Configuring the iSCSI Proxy Initiator	20-16
Access Control in iSCSI	20-16
Fibre Channel Zoning-Based Access Control	20-16
iSCSI-Based Access Control	20-17
Enforcing Access Control	20-17
iSCSI User Authentication	20-17
No Authentication	20-18
Configuring an Authentication Mechanism	20-18
Configuring an iSCSI RADIUS Server	20-18
Advanced iSCSI Configuration	20-19
Setting the QoS Values	20-19
iSCSI Forwarding Mode	20-20
iSCSI High Availability	20-20
Configuring Storage Name Services	20-23
Creating iSNS Profiles and Tagging Profiles	20-23
Creating an iSNS Profile	20-24
Modifying an iSNS Profile	20-24

## CHAPTER 21

## FICON Configuration 21-1

About FICON	21-2
MDS-Specific FICON Advantages	21-2
Fabric-Optimization with VSANs	21-2
FCIP Support	21-3
PortChannel Support	21-4
VSANs for FICON and FCP Intermixing	21-4
Cisco MDS-Supported FICON Features	21-4
FICON Port Numbering	21-6
FICON Port Numbering Guidelines	21-7
FCIP and PortChannel Port Numbers	21-8
Port Addresses	21-8
Installed and Uninstalled Ports	21-8
FC ID Allocation	21-8
FICON Cascading	21-9
FICON VSAN Prerequisites	21-9
Enabling FICON	21-10
Creating FICON VSANs and enabling FICON	21-10

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Deleting FICON VSANs	21-11
Viewing FICON Director History	21-12
The code-page Option	21-12
FC ID Last Byte	21-12
FICON Host Control	21-13
Host Changes FICON Port Parameters	21-13
FICON Information Refresh Note	21-14
Configuring FICON Ports	21-14
Port Blocking	21-14
Port Prohibiting	21-14
Configuring Port Blocking and Port Prohibiting	21-15
Entering FICON Port Configuration Information	21-15
Viewing FICON Port Attributes	21-16
FICON Configuration Files	21-16
Accessing FICON Configuration Files	21-17
Copying FICON Configuration Files	21-17
Editing FICON Configuration Files	21-17
Managing FICON Configuration Files In Device Manager	21-18
Port Swapping	21-18
Port Swapping Guidelines	21-19
Swapping FICON Ports	21-19
Clearing FICON Device Allegiance	21-19
CUP In-Band Management	21-20
Fabric Binding Configuration	21-20
Port Security Versus Fabric Binding	21-20
Fabric Binding Enforcement	21-21
Enabling Fabric Binding	21-21
Configuring a List of Switch WWNs In a Fabric	21-22
Activating Fabric Binding	21-22
Saving Fabric Binding Configurations	21-23
Deactivating Fabric Binding	21-23
Fabric Binding CopyActive to Config	21-23
Creating a Fabric Binding Configuration	21-24
Deleting a Fabric Binding Configuration	21-24
Viewing Fabric Binding Active Database	21-24
Viewing Fabric Binding Violations	21-24
Clearing Fabric Binding Statistics	21-25
Viewing EFMD Statistics	21-25
Displaying RLIR Information	21-25

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Calculating FICON Flow Load Balance 21-25

## CHAPTER 22

### Configuring Intelligent Storage Services 22-1

Intelligent Storage Services 22-1

Enabling Intelligent Storage Services 22-2

Disabling Intelligent Storage Services 22-3

SCSI Flow Services 22-3

Configuring SCSI Flow Services 22-4

Fibre Channel Write Acceleration 22-4

Configuring Fibre Channel Write Acceleration 22-5

SCSI Flow Statistics 22-5

Enabling SCSI Flow Statistics 22-6

Viewing SCSI Flow Statistics and Clearing SCSI Flow Statistics 22-6

SANTap 22-7

Transparent Mode 22-8

Proxy Mode-1 22-9

Proxy Mode-2 22-10

Configuring SANTap 22-10

NASB 22-11

Configuring NASB 22-12

## CHAPTER 23

### Additional Configuration 23-1

Fibre Channel Time Out Values 23-1

The fctrace Feature 23-2

Performing an fctrace Operation 23-2

The fcping Feature 23-2

Invoking the fcping Feature 23-3

Configuring World Wide Names 23-3

Link Initialization WWN Usage 23-3

Flat FC ID Allocation 23-4

Loop Monitoring Initiation 23-4

Switch Interoperability 23-4

Interoperability Configuration 23-6

Configuring Interoperability 23-6

## PART 4

### Security Configuration

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

---

**CHAPTER 24**
**Users and Common Roles 24-1**

- Role-Based Authorization 24-1
- Configuring Common Roles 24-2
  - Creating Common Roles 24-2
  - Editing Rules For Common Roles in Device Manager 24-3
  - Deleting Common Roles 24-3
  - Configuring the VSAN Policy 24-3
    - Modifying the VSAN Policy 24-4
- Configuring User Accounts 24-4
  - Creating or Updating Users 24-5
    - Creating Strong Passwords 24-5
  - Adding a User 24-5
  - Deleting a User 24-6
  - Viewing User Information 24-6
- Configuring SSH Services 24-6
  - Generating the SSH Server Key Pair and Enabling SSH 24-6
  - Deleting a Generated Key Pair 24-7
  - Recovering Administrator Password 24-7

---

**CHAPTER 25**
**SNMP Configuration 25-1**

- About SNMP 25-1
  - SNMP Version 1 and Version 2c 25-2
  - SNMP Version 3 25-2
  - SNMP v3 CLI User Management and AAA Integration 25-2
    - CLI and SNMP User Synchronization 25-2
    - Software Upgrade Synchronization 25-3
  - Restricting Switch Access 25-3
  - Adding a Community String 25-3
  - Deleting a Community String 25-3
- Understanding Users 25-4
  - Adding a User 25-4
  - Deleting a User 25-4
  - Viewing SNMP Community and User Information 25-5
  - Group-Based SNMP Access 25-5
- Assigning SNMPv3 Users to Multiple Roles 25-5
- Configuring SNMP Notifications 25-6



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## **CHAPTER 26**

### **RADIUS and TACACS+ 26-1**

Authentication, Authorization, and Accounting	26-1
CLI Security Options	26-1
SNMP Security Options	26-2
Switch AAA Functionalities	26-2
Authentication	26-2
Authorization	26-2
Accounting	26-3
Remote AAA Services	26-4
Remote Authentication Guidelines	26-4
Server Groups	26-4
AAA Service Configuration Options	26-4
Configuring RADIUS	26-5
Setting the RADIUS Server for Authentication and Accounting	26-5
Setting the Global Preshared Key	26-6
Defining Vendor-Specific Attributes	26-6
VSA Format	26-6
Specifying SNMPv3 on AAA Servers	26-7
Configuring TACACS+	26-7
About TACACS+	26-8
Enabling TACACS+	26-8
Setting the TACACS+ Server	26-8
Defining Custom Attributes for Roles	26-8
Supported TACACS+ Servers	26-9
Configuring Server Groups	26-9
Distributing AAA server Configuration	26-10
Enabling the distribution	26-10
Starting a Distribution Session on a Switch	26-10
Committing the Distribution	26-11
Discarding the Distribution Session	26-11
Local AAA Services	26-11
Disabling AAA Authentication	26-11

## **CHAPTER 27**

### **IP Access Control Lists 27-1**

IP-ACL Configuration Guidelines	27-1
Filter Contents	27-2
Protocol Information	27-2
Address Information	27-2
Port Information	27-3

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Using the IP-ACL Wizard	27-4
Creating Complex IP-ACLs Using Device Manager	27-5
Associating IP-ACL Profiles to Interfaces	27-6
Removing Associations Between IP-ACL Profiles and Interfaces	27-6
Deleting IP Profiles	27-7

## CHAPTER 28

### **IPsec and IKE 28-1**

Configuring IPsec Network Security	28-1
The 14/2-Port Multiprotocol Services Module	28-1
IPsec Prerequisites	28-2
IPsec Compatibility	28-3
About IPsec	28-3
About IKE	28-4
IPsec and IKE Terminology	28-4
Supported IPsec Transforms	28-5
Supported IKE Transforms and Algorithms	28-6
Supported Algorithms for Windows and Linux Platforms	28-7
Enabling IPsec Using FCIP Wizard	28-7
Modifying IKE and IPsec	28-8
Crypto ACL Guidelines	28-9
Mirror Image Crypto ACLs	28-10
The any Keyword in Crypto ACLs	28-12
Configuring Crypto IP-ACLs	28-12
Transform Sets	28-12
Crypto Map Entries	28-13
SA Establishment Between Peers	28-14
The AutoPeer Option	28-14
SA Lifetime Negotiation	28-15
Perfect Forwarding Secrecy	28-15
Creating or Modifying Crypto Maps	28-15
Applying a Crypto Map Set to an Interface	28-16
IPsec Maintenance	28-17
Global Lifetime Values	28-17

## CHAPTER 29

### **FC-SP and DHCHAP 29-1**

Fibre Channel Security Protocol	29-1
About DHCHAP	29-2
DHCHAP Compatibility with Existing Cisco MDS Features	29-3

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Configuring DHCPAP Authentication	29-3
Enabling DHCPAP	29-3
Configuring DHCPAP Authentication Modes	29-4
Changing the DHCPAP Hash Algorithm	29-5
Changing DHCPAP Group Settings	29-5
Configuring the DHCPAP Password	29-6
Configuring the DHCPAP Password for the Local Switch	29-7
Configuring Remote Passwords for Other Devices	29-7
Setting the DHCPAP Timeout Value	29-8
Configuring DHCPAP AAA Authentication	29-8
Enabling FC-SP on ISLs	29-8

## CHAPTER 30

### Port Security 30-1

About Port Security	30-1
About Auto-Learn	30-1
Auto-Learning Device Authorization	30-2
Port Security Enforcement	30-2
Configuring Port Security	30-3
Enabling Port Security	30-3
Activating Port Security with Auto-Learn	30-3
Displaying Activated Port Security Settings	30-4
Displaying Port Security Statistics	30-4
Displaying Port Security Violations	30-4
Turning Auto-Learning On or Off	30-5
Example of Port Security Authorization	30-5
Port Security Manual Configuration	30-6
WWN Identification	30-6
Manually Configuring Port Security	30-7
Deleting a Port Security Pair	30-7
Database Interaction	30-8
Database Scenarios	30-9
Activating the Port Security Database	30-10
Database Activation Rejection	30-10
Forceful Port Security Activation	30-11
Database Reactivation	30-11
Copying an Active Database to the Config Database	30-11

## PART 5

### Network and Performance Monitoring

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

---

**CHAPTER 31**
**Network Monitoring 31-1**

- SAN Discovery and Topology Mapping 31-1
  - Device Discovery 31-1
  - Topology Mapping 31-1
    - Using the Topology Map 31-2
    - Saving a Customized Topology Map Layout 31-2
    - Using Enclosures with Fabric Manager Topology Maps 31-2
    - Mapping Multiple Fabrics 31-3
  - Inventory Management 31-3
    - Using the Inventory Tab from Fabric Manager Web Services 31-3
- Configuring System Message Logging 31-4
  - Syslog Server Logging Facilities and Severity Levels 31-4
  - Configuring Message Logging 31-7
  - Configuring a Syslog Server 31-8
  - Verifying Syslog Servers from Fabric Manager Web Services 31-9
  - Viewing Logs from Fabric Manager Web Services 31-9
  - Viewing Logs from Device Manager 31-9
- Health and Event Monitoring 31-9
  - Fabric Manager Events Tab 31-10
  - Event Information in Fabric Manager Web Services Reports 31-10
  - Events in Device Manager 31-10

---

**CHAPTER 32**
**Performance Monitoring 32-1**

- Real-Time Performance Monitoring 32-1
  - Device Manager Real-Time Performance Monitoring 32-1
  - Fabric Manager Real-Time ISL Statistics 32-2
- Historical Performance Monitoring 32-2
  - Creating a Flow with Performance Manager 32-3
  - Creating a Collection with Performance Manager 32-3
    - Using Performance Thresholds 32-4
    - Using the Performance Manager Configuration Wizard 32-5
  - Starting and Stopping Data Collection 32-6
  - Viewing Performance Manager Reports 32-6
    - Performance Summary 32-6
    - Performance Tables and Details Graphs 32-7
    - Viewing Performance Manager Events 32-7
  - Generating Top10 Reports in Performance Manager 32-7
    - Generating Top10 Reports Using Scripts 32-8
  - Exporting Data Collections to XML Files 32-8

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Exporting Data Collections in Readable Format 32-9

Configuring Performance Manager for Use with Cisco Traffic Analyzer 32-10

## CHAPTER 33

### Third-Party Integration 33-1

Call Home Configuration 33-1

Cisco AutoNotify 33-1

Configuring Call Home 33-2

Configuring Call Home Destination Profiles and Alert Groups 33-2

Call Home Message Severity Levels 33-3

Event Triggers 33-3

Message Contents 33-5

Configuring SNMP Events 33-11

Filtering SNMP Events 33-11

Configuring SNMP Event Destinations 33-12

Configuring Event Security 33-12

Viewing the SNMP Events Log 33-13

Configuring RMON Using Threshold Manager 33-13

Enabling RMON Alarms by Port 33-13

Enabling RMON Alarms for VSANs 33-14

Enabling RMON Alarms for Physical Components 33-14

Managing RMON Events 33-15

Managing RMON Alarms 33-15

Viewing the RMON Log 33-16

## PART 6

### Network Troubleshooting

## CHAPTER 34

### Troubleshooting Your Fabric 34-1

Troubleshooting Tools and Techniques 34-1

Cisco Traffic Analyzer 34-2

Cisco Protocol Analyzer 34-3

Analyzing Switch Device Health 34-3

Online System Health Management 34-4

Loopback Test Configuration Frequency 34-4

Performing Internal Loopbacks 34-4

Performing External Loopbacks 34-5

Hardware Failure Action 34-5

Analyzing Switch Fabric Configuration 34-5

Analyzing End-to-End Connectivity 34-6

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Configuring a Fabric Analyzer	34-7
About the Cisco Fabric Analyzer	34-7
Local Text-Based Capture	34-8
Remote Capture Daemon	34-8
GUI-Based Client	34-9
Configuring the Cisco Fabric Analyzer	34-9
Sending Captures to Remote IP Addresses	34-9
Displaying Captured Frames	34-10
Defining Display Filters	34-11
Capture Filters	34-11
Permitted Capture Filters	34-12
Using the Ping Tool	34-12
Using Traceroute and Other Troubleshooting Tools	34-13
Analyzing the Results of Merging Zones	34-13
Issuing the Show Tech Support Command	34-14
Locating Other Switches	34-15
Getting Oversubscription Information in Device Manager	34-16

## CHAPTER 35

### Management Software Troubleshooting 35-1

Installation Issues	35-3
When installing Fabric Manager from windows, clicking the install button fails.	35-3
How do I install Java Web Start on a UNIX machine?	35-4
Why can't I launch Fabric Manager on Solaris?	35-4
Why is my browser prompting to save JNLP files?	35-4
Why do I get a "Java Web Start not detected" error?	35-4
Why can't I see my desktop shortcuts?	35-5
How do I upgrade to a newer version?	35-5
How do I downgrade Fabric Manager or Device Manager?	35-5
What do I do if my upgrade is not working?	35-5
Java Web Start hangs on download dialog. What do I do?	35-5
How can I manually configure my browser for Java Web Start?	35-6
Can I run Java Web Start from the command line?	35-6
Windows 2000 crashes (blue screen). What do I do?	35-6
How do I clear the Java Web Start cache?	35-7
Why doesn't my login work in Fabric Manager and Device Manager?	35-7
Why can't I install Fabric Manager or Device Manager when pcAnywhere is running?	35-7
The Fabric Manager or the Performance Manager service shows up as "disabled" in the Services menu.	35-7

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Why can't I install Fabric Manager or Device Manager when McAfee Internet Suite 6.0 Professional is running? **35-8**

I get an error ".sm/logon." when I downgrade from MDS SAN-OS Release 2.x (or newer) to 1.3(x). **35-8**

#### **General 35-8**

Why do I get errors while monitoring Area chart graphing? **35-8**

Why do I get "gen error" messages? **35-8**

Why are disk images in the Device Manager Summary View not showing up? **35-8**

Why can't I set both the D\_S\_TOV and E\_D\_TOV timers in the Device Manager? **35-9**

Why are the columns in the Device Manager tables too small? **35-9**

Why are my fabric changes not propagated onto the map (for example, links don't disappear)? **35-9**

Why does the PortChannel creation dialog become too small after several uses? **35-9**

Why do I see errors when I have configured IPFC? **35-9**

Why is Fabric Manager or Device Manager using the wrong network interface? **35-9**

Why am I seeing display anomalies? **35-10**

How do I connect the Fabric Manager client to the server across VPN? **35-10**

Why is the active zone set in edit zone always shown in bold (even after successful activation)? **35-10**

Can I create a zone with prefix IVRZ and a zone set with name nozonset? **35-10**

One-click license install fails, cannot connect to Cisco website. **35-10**

Fabric Manager client and Device Manager cannot connect to the switch **35-10**

License Wizard fails to fetch license keys, saying connect failed **35-11**

How do I increase log window size in Fabric Manager Client? **35-11**

#### **Windows Issues 35-11**

Text fields showing up too small, cannot enter any data **35-11**

Why does CiscoWorks fail to start in the browser? **35-11**

Help contents are unreadable because of highlighting **35-11**

Printing causes an application crash **35-11**

Windows XP hangs (or blue screen). What do I do? **35-12**

Why do the Device Manager Icons Disappear Sometimes? **35-12**

Why does Fabric Manager hang when I drag an existing Zone Member to a Zone? **35-12**

Device Manager or Fabric Manager window content disappears in Windows XP **35-12**

Why does SCP/SFTP fail when I try to copy a file from my local machine to the switch? **35-12**

#### **UNIX Issues 35-13**

Why Do the Parent Menus Disappear? **35-13**

Why do I keep getting a "too many open files" error? **35-13**

#### **Other 35-14**

How can I set the map layout so it stays after I restart Fabric Manager? **35-14**

Two switches show on my map, but I only have one switch **35-14**

There is a red/orange/dotted line through the switch. What's wrong? **35-14**

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Can I upgrade without losing my map settings?	35-19
How can I preserve historical data when moving Fabric Manager server to a new host?	35-20
Are There Any Restrictions When Using Fabric Manager Across FCIP?	35-20
I see "Please insure that FM server is running on localhost."	35-20
How can I run Cisco Fabric Manager if I have multiple interfaces?	35-21
Manually specifying an interface for Fabric Manager Server	35-21
Manually specifying an interface for Fabric Manager Client or Device Manager	35-22
How can I configure an HTTP proxy server?	35-22
How can I clear the topology map?	35-23
Can I use Fabric Manager in a mixed software environment?	35-23
I Get an Error When Launching Fabric Manager	35-23
Can I Search for Devices in a Fabric?	35-24
Do I Need A License of Fabric Manager Server for Each Switch in the Fabric?	35-24
How can I Manage Multiple Fabrics?	35-24
License Expiration Causes Orange X Through Switch	35-24

---

**APPENDIX A**

**GUI/CLI Usage Chart A-1**

Procedures A-1

---

**APPENDIX B**

**Interface Nonoperational Reason Codes B-1**

---

**INDEX**



## New and Changed Information

The *Cisco MDS 9000 Family Fabric Manager Configuration Guide* provides release-specific information for the Cisco MDS SAN-OS Release 2.x (including Release 2.0(1b) through 2.1(1a) software). The latest version of this document is available at the following Cisco Systems website: <http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm>

To check for additional information about this release, refer to the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website: <http://www.cisco.com/univercd/cc/td/doc/product/sn5000/mds9000/index.htm>

**Table 1** summarizes the new and changed features for the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*, and tells you where they are documented. The table includes a brief description of each new feature and the release in which the change occurred.

**Table 1**      **Documented Features for the Cisco MDS 9000 Family Fabric Manager Configuration Guide**

Feature	Description	Changed in Release	Where Documented
Network-Accelerated Serverless Backup	Offloads data movement to Network-Accelerated Serverless Backup (NASB) devices that use SCSI Xcopy.	2.1(1a)	<a href="#">Chapter 22, “Configuring Intelligent Storage Services”</a>
SANTap	Integrates third-party data storage application into the SAN.	2.1(1a)	<a href="#">Chapter 22, “Configuring Intelligent Storage Services”</a>
IVR Enhancements	Supports IVR NAT, IVR auto-topology, and autonomous fabric IDs.	2.1(1a)	<a href="#">Chapter 16, “Inter-VSAN Routing Configuration”</a>
Performance Manager Top 10 Reports	Generates historical performance reports based on the top 10 entities in the report category.	2.1(1a)	<a href="#">Chapter 32, “Performance Monitoring”</a>
Fibre Channel Write Acceleration	Minimizes application latency or reduces transactions per second over long distances (supported by Fabric Manager).	2.0(2b)	<a href="#">Chapter 22, “Configuring Intelligent Storage Services”</a>
SCSI Flow Statistics	Collects statistics for configured SCSI flows (supported by Fabric Manager).	2.0(2b)	<a href="#">Chapter 22, “Configuring Intelligent Storage Services”</a>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 1**      ***Documented Features for the Cisco MDS 9000 Family Fabric Manager Configuration Guide (continued)***

<b>Feature</b>	<b>Description</b>	<b>Changed in Release</b>	<b>Where Documented</b>
New Module	Storage Services Module or SSM (supported by Device Manager).	2.0(2b)	<a href="#">Chapter 4, “Device Manager”</a>
Fabric Manager Enhancements	Exports Performance Manager reports in comma-separated format.  Generates Top10 Reports in Performance Manager.	2.1(1a)	<a href="#">Chapter 32, “Performance Monitoring”</a>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 1 Documented Features for the Cisco MDS 9000 Family Fabric Manager Configuration Guide (continued)**

Feature	Description	Changed in Release	Where Documented
Fabric Manager Enhancements	<p>Displays SANs and multiple fabrics in Fabric pane.</p> <p>Filtering.</p> <p>Rearranged Logical and Physical panes.</p> <p>Detachable tables in Information pane.</p> <p>Persist fabrics for monitoring by Fabric Manager Server.</p> <p>Login screen enhancements include simple versus complex displays, ability to load from the database, ability to sync server to same NIC as client.</p> <p>Enclosures in map bring up customized app when right-clicked.</p> <p>Displays every VSAN island without collapsing.</p> <p>Enhanced zoning capabilities.</p> <p>AES Support (authentication algorithm).</p> <p>SCSI target IDs are now associated with targets.</p> <p>FDMI and Name Server information is collated for initiators (hosts).</p> <p>Enclosures are global across SANs.</p> <p>FCIP Wizard enhancements include encryption and compression.</p> <p>FICON enhancements include ability to display FICON port numbers on map, and ability to assign FICON ports for FCIP PortChannels.</p> <p>Zoning enhancements include aliases treated as groups; multiple alias types; ability to rename zone sets, zones, and aliases; backup and restore zone database; and enhanced zoning.</p> <p>Release 2.0(1b) feature support, including DPVM Wizard, Cisco Fabric Services, zone-based QoS, IKE/IPsec, port tracking, and DNS.</p>	2.0(1b)	<a href="#">Chapter 3, “Fabric Manager Client”</a>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 1**      ***Documented Features for the Cisco MDS 9000 Family Fabric Manager Configuration Guide (continued)***

<b>Feature</b>	<b>Description</b>	<b>Changed in Release</b>	<b>Where Documented</b>
Device Manager Enhancements	<p>MPS 14/2 support.</p> <p>AES support (authentication algorithm).</p> <p>FCIP interfaces displayed in Physical View.</p> <p>Release 2.0(1b) feature support, including auto-trunk, port tracking, DNS, tape acceleration, IPS encryption, Cisco Fabric Services, and DPVM.</p> <p>Gigabit Ethernet TCP statistics.</p> <p>Multicast root.</p> <p>FCID area allocation.</p> <p>Additional (and more accurate) Flash file manipulation capabilities.</p> <p>Ability to read syslog information from FM Server.</p> <p>Summary View enhancements including display of EtherChannel members, which Gigabit Ethernet port is associated with FCIP, and FCIP compression information.</p> <p>Ability to power down a line card.</p>	2.0(1b)	<a href="#">Chapter 4, “Device Manager”</a>
Performance Manager Enhancements	<p>Summary and drill down report; Quick Data Collector and Flow Setup Wizards; wizard enhancements include interpolation, adaptive baseline thresholds, compression, and enhanced collection capabilities</p>	2.0(1b)	<a href="#">Chapter 31, “Network Monitoring”</a>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 1**      ***Documented Features for the Cisco MDS 9000 Family Fabric Manager Configuration Guide (continued)***

<b>Feature</b>	<b>Description</b>	<b>Changed in Release</b>	<b>Where Documented</b>
Fabric Manager Web Services	Accesses Performance Manager data. Information (events, performance data, inventory summary) can be viewed remotely, is displayed quickly, and can be filtered; baseline thresholds can be defined and triggered; custom periods can be defined (in addition to day/week/month/year); also has some server administrative capabilities.	2.0(1b)	<a href="#">Chapter 5, “Fabric Manager Web Services”</a>
Fabric Manager Server Enhancements	Multiple fabric management, centralized discovery, continuous health monitoring, roaming user profiles	2.0(1b)	<a href="#">Chapter 2, “Fabric Manager Server”</a>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. It also provides information on how to obtain related documentation.

## Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

You should be familiar with the basic concepts and terminology used in internetworking, and understand your network topology and the protocols that the devices in your network can use. You should also have a working knowledge of the operating system on which you are running Fabric Manager, such as Microsoft Windows, Linux, or Solaris.

## Organization

This guide describes the most commonly used features of Fabric Manager and Device Manager. Refer to the online help available with Fabric Manager or Device Manager for details on all features.

This guide is organized as follows:

Chapter	Title	Description
Chapter 1	<a href="#">Installation and Configuration</a>	Provides a brief overview of Fabric Manager components and capabilities, and information on installation and launching the applications.
Chapter 2	<a href="#">Fabric Manager Server</a>	Provides in-depth descriptions of GUIs and capabilities for the Fabric Manager Server.
Chapter 3	<a href="#">Fabric Manager Client</a>	Provides in-depth descriptions of GUIs and capabilities for the Fabric Manager.
Chapter 4	<a href="#">Device Manager</a>	Provides in-depth descriptions of GUIs and capabilities for the Device Manager.
Chapter 5	<a href="#">Fabric Manager Web Services</a>	Provides in-depth descriptions of GUIs and capabilities for the Fabric Manager Web Client.
Chapter 6	<a href="#">Performance Manager</a>	Provides overview of Performance Manager architecture.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Chapter	Title	Description
Chapter 7	<a href="#">Authentication in Fabric Manager</a>	Describes the authentication schemes between Fabric Manager components and fabric switches.
Chapter 8	<a href="#">Cisco Traffic Analyzer</a>	Describes installing and launching Cisco Traffic Analyzer from Performance Manager.
Chapter 9	<a href="#">Obtaining and Installing Licenses</a>	Provides information on the Cisco MDS 9000 Family licensing model, license concepts, and license installation and management.
Chapter 10	<a href="#">Software Images</a>	Describes how to upgrade Cisco MDS 9000 Family switches, install software image files, use the Flash file system on the supervisor engine, and recover a corrupted bootflash image.
Chapter 11	<a href="#">Configuration Files</a>	Describes how to update configuration files.
Chapter 12	<a href="#">Cisco Fabric Services</a>	Describes Cisco Fabric Services, used for distributing configuration changes through the fabric.
Chapter 13	<a href="#">VSAN Configuration</a>	Describes how virtual SANs (VSANs) work, explains the concept of default VSANs, isolated VSANs, VSAN IDs, and attributes, and provides details on how to create, delete, and view VSANs.
Chapter 14	<a href="#">Dynamic VSAN Configuration</a>	Describes how to dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature.
Chapter 15	<a href="#">Zone Configuration</a>	Defines various zoning concepts and provides details on configuring a zone set and zone management features.
Chapter 16	<a href="#">Inter-VSAN Routing Configuration</a>	Provides details on sharing resources across VSANs using the inter-VSAN Routing (IVR) feature
Chapter 17	<a href="#">PortChannel Configuration</a>	Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.
Chapter 18	<a href="#">Interface Configuration</a>	Explains port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
Chapter 19	<a href="#">FCIP Configuration</a>	Provides details on extending the reach of Fibre Channel SANs by connecting separated SAN islands together through IP networks using FCIP.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Chapter	Title	Description
Chapter 20	<a href="#">iSCSI Configuration</a>	Provides details on extending the reach of Fibre Channel SANs by allowing IP hosts to access FC storage using the iSCSI protocol.
Chapter 21	<a href="#">FICON Configuration</a>	Provides details on the FI-bre CON-nection (FICON) interface, fabric binding, and the Registered Link Incident Report (RLIR) capabilities in Cisco MDS switches.
Chapter 22	<a href="#">Configuring Intelligent Storage Services</a>	Describes the intelligent storage services available on the Storage Services Module (SSM), including Fibre Channel write acceleration and SCSI flow statistics.
Chapter 23	<a href="#">Additional Configuration</a>	Describes the advanced configuration features—time out values, fctrace, fabric analyzer, world wide names, flat FC IDs, loop monitoring, and interoperating switches.
Chapter 24	<a href="#">Users and Common Roles</a>	Describes Common user roles and SSH.
Chapter 25	<a href="#">SNMP Configuration</a>	Describes SNMP security, notifications, and user roles.
Chapter 26	<a href="#">RADIUS and TACACS+</a>	Describes RADIUS and TACACS+ authorization and accounting services.
Chapter 27	<a href="#">IPsec and IKE</a>	Describes IPsec, and configuration through Fabric Manager.
Chapter 28	<a href="#">FC-SP and DHCHAP</a>	Describes Fibre Channel Security Protocol and how to configure DHCHAP to work with FCSP.
Chapter 29	<a href="#">IP Access Control Lists</a>	Describes controlling network access through IP ACLs.
Chapter 30	<a href="#">Port Security</a>	Describes how to control access to the fabric through port security.
Chapter 31	<a href="#">Network Monitoring</a>	Describes SAN topology, inventory, and event monitoring.
Chapter 32	<a href="#">Performance Monitoring</a>	Describes real-time and historical performance monitoring using Fabric Manager and Performance Manager.
Chapter 33	<a href="#">Third-Party Integration</a>	Describes integrating SNMP, syslog, and Call Home with third party management applications.
Chapter 34	<a href="#">Troubleshooting Your Fabric</a>	Provides information on using Fabric Manager to troubleshoot your fabric.
Chapter 35	<a href="#">Management Software Troubleshooting</a>	Answers some of the most frequently asked questions about Cisco Fabric Manager.
Appendix A	<a href="#">GUI/CLI Usage Chart</a>	Provides a table of procedures, organized by best performed by the CLI, Fabric Manager, or Device Manager.
Appendix B	<a href="#">Interface Nonoperational Reason Codes</a>	Explains the reason codes for why an interface is operationally down.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Document Conventions

Command descriptions use these conventions:

<b>boldface font</b>	Commands and keywords are in boldface.
<i>italic font</i>	Arguments for which you supply values are in italics.

Screen examples use these conventions:

<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<b>boldface screen font</b>	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
< >	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:



### Note

Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.



### Caution

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

## Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents:

- *Cisco MDS 9000 Family Release Notes for Cisco MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Interoperability Support Matrix*
- *Cisco MDS SAN-OS Release Compatibility Matrix for IBM SAN Volume Controller Software for Cisco MDS 9000*
- *Cisco MDS SAN-OS Release Compatibility Matrix for VERITAS Storage Foundation for Networks Software*
- *Cisco MDS SAN-OS Release Compatibility Matrix for SSI Images*
- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9000 Family Software Upgrade and Downgrade Guide*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Fabric Manager Configuration Guide*
- *Cisco MDS 9000 Family Fabric and Device Manager Online Help*
- *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide*
- *Cisco MDS 9000 Family MIB Quick Reference*
- *Cisco MDS 9000 Family CIM Programming Reference*
- *Cisco MDS 9000 Family System Messages Reference*
- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*
- *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Note*

For information on VERITAS Storage Foundation™ for Networks for the Cisco MDS 9000 Family, refer to the VERITAS website: <http://support.veritas.com/>

For information on IBM TotalStorage SAN Volume Controller Storage Software for the Cisco MDS 9000 Family, refer to the IBM TotalStorage Support website:  
<http://www.ibm.com/storage/support/2062-2300/>

## Obtaining Documentation

Cisco documentation and additional literature are available on Cisco.com. Cisco also provides several ways to obtain technical assistance and other technical resources. These sections explain how to obtain technical information from Cisco Systems.

### Cisco.com

You can access the most current Cisco documentation at this URL:

<http://www.cisco.com/univercd/home/home.htm>

You can access the Cisco website at this URL:

<http://www.cisco.com>

You can access international Cisco websites at this URL:

[http://www.cisco.com/public/countries\\_languages.shtml](http://www.cisco.com/public/countries_languages.shtml)

### Documentation DVD

Cisco documentation and additional literature are available in a Documentation DVD package, which may have shipped with your product. The Documentation DVD is updated regularly and may be more current than printed documentation. The Documentation DVD package is available as a single unit.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Registered Cisco.com users (Cisco direct customers) can order a Cisco Documentation DVD (product number DOC-DOCDVD=) from the Ordering tool or Cisco Marketplace.

Cisco Ordering tool:

<http://www.cisco.com/en/US/partner/ordering/>

Cisco Marketplace:

<http://www.cisco.com/go/marketplace/>

## Ordering Documentation

You can find instructions for ordering documentation at this URL:

[http://www.cisco.com/univercd/cc/td/doc/es\\_inpk/pdi.htm](http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm)

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Ordering tool:  
<http://www.cisco.com/en/US/partner/ordering/>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408 526-7208 or, elsewhere in North America, by calling 1 800 553-NETS (6387).

## Documentation Feedback

You can send comments about technical documentation to [bug-doc@cisco.com](mailto:bug-doc@cisco.com).

You can submit comments by using the response card (if present) behind the front cover of your document or by writing to the following address:

Cisco Systems  
Attn: Customer Document Ordering  
170 West Tasman Drive  
San Jose, CA 95134-9883

We appreciate your comments.

## Cisco Product Security Overview

Cisco provides a free online Security Vulnerability Policy portal at this URL:

[http://www.cisco.com/en/US/products/products\\_security\\_vulnerability\\_policy.html](http://www.cisco.com/en/US/products/products_security_vulnerability_policy.html)

From this site, you can perform these tasks:

- Report security vulnerabilities in Cisco products.
- Obtain assistance with security incidents that involve Cisco products.
- Register to receive security information from Cisco.

A current list of security advisories and notices for Cisco products is available at this URL:

<http://www.cisco.com/go/psirt>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

If you prefer to see advisories and notices as they are updated in real time, you can access a Product Security Incident Response Team Really Simple Syndication (PSIRT RSS) feed from this URL:

[http://www.cisco.com/en/US/products/products\\_psirt\\_rss\\_feed.html](http://www.cisco.com/en/US/products/products_psirt_rss_feed.html)

## Reporting Security Problems in Cisco Products

Cisco is committed to delivering secure products. We test our products internally before we release them, and we strive to correct all vulnerabilities quickly. If you think that you might have identified a vulnerability in a Cisco product, contact PSIRT:

- Emergencies—[security-alert@cisco.com](mailto:security-alert@cisco.com)
- Nonemergencies—[psirt@cisco.com](mailto:psirt@cisco.com)



We encourage you to use Pretty Good Privacy (PGP) or a compatible product to encrypt any sensitive information that you send to Cisco. PSIRT can work from encrypted information that is compatible with PGP versions 2.x through 8.x.

Never use a revoked or an expired encryption key. The correct public key to use in your correspondence with PSIRT is the one that has the most recent creation date in this public key server list:

<http://pgp.mit.edu:11371/pks/lookup?search=psirt%40cisco.com&op=index&exact=on>

In an emergency, you can also reach PSIRT by telephone:

- 1 877 228-7302
- 1 408 525-6532

## Obtaining Technical Assistance

For all customers, partners, resellers, and distributors who hold valid Cisco service contracts, Cisco Technical Support provides 24-hour-a-day, award-winning technical assistance. The Cisco Technical Support Website on Cisco.com features extensive online support resources. In addition, Cisco Technical Assistance Center (TAC) engineers provide telephone support. If you do not hold a valid Cisco service contract, contact your reseller.

## Cisco Technical Support Website

The Cisco Technical Support Website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies. The website is available 24 hours a day, 365 days a year, at this URL:

<http://www.cisco.com/techsupport>

Access to all tools on the Cisco Technical Support Website requires a Cisco.com user ID and password. If you have a valid service contract but do not have a user ID or password, you can register at this URL:

<http://tools.cisco.com/RPF/register/register.do>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

Use the Cisco Product Identification (CPI) tool to locate your product serial number before submitting a web or phone request for service. You can access the CPI tool from the Cisco Technical Support Website by clicking the **Tools & Resources** link under Documentation & Tools. Choose **Cisco Product Identification Tool** from the Alphabetical Index drop-down list, or click the **Cisco Product Identification Tool** link under Alerts & RMAs. The CPI tool offers three search options: by product ID or model name; by tree view; or for certain products, by copying and pasting **show** command output. Search results show an illustration of your product with the serial number label location highlighted. Locate the serial number label on your product and record the information before placing a service call.

## Submitting a Service Request

Using the online TAC Service Request Tool is the fastest way to open S3 and S4 service requests. (S3 and S4 service requests are those in which your network is minimally impaired or for which you require product information.) After you describe your situation, the TAC Service Request Tool provides recommended solutions. If your issue is not resolved using the recommended resources, your service request is assigned to a Cisco TAC engineer. The TAC Service Request Tool is located at this URL:

<http://www.cisco.com/techsupport/servicerequest>

For S1 or S2 service requests or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 service requests are those in which your production network is down or severely degraded.) Cisco TAC engineers are assigned immediately to S1 and S2 service requests to help keep your business operations running smoothly.

To open a service request by telephone, use one of the following numbers:

Asia-Pacific: +61 2 8446 7411 (Australia: 1 800 805 227)

EMEA: +32 2 704 55 55

USA: 1 800 553-2447

For a complete list of Cisco TAC contacts, go to this URL:

<http://www.cisco.com/techsupport/contacts>

## Definitions of Service Request Severity

To ensure that all service requests are reported in a standard format, Cisco has established severity definitions.

**Severity 1 (S1)**—Your network is “down,” or there is a critical impact to your business operations. You and Cisco will commit all necessary resources around the clock to resolve the situation.

**Severity 2 (S2)**—Operation of an existing network is severely degraded, or significant aspects of your business operation are negatively affected by inadequate performance of Cisco products. You and Cisco will commit full-time resources during normal business hours to resolve the situation.

**Severity 3 (S3)**—Operational performance of your network is impaired, but most business operations remain functional. You and Cisco will commit resources during normal business hours to restore service to satisfactory levels.

**Severity 4 (S4)**—You require information or assistance with Cisco product capabilities, installation, or configuration. There is little or no effect on your business operations.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Obtaining Additional Publications and Information

Information about Cisco products, technologies, and network solutions is available from various online and printed sources.

- Cisco Marketplace provides a variety of Cisco books, reference guides, and logo merchandise. Visit Cisco Marketplace, the company store, at this URL:

<http://www.cisco.com/go/marketplace/>

- *Cisco Press* publishes a wide range of general networking, training and certification titles. Both new and experienced users will benefit from these publications. For current Cisco Press titles and other information, go to Cisco Press at this URL:

<http://www.ciscopress.com>

- *Packet* magazine is the Cisco Systems technical user magazine for maximizing Internet and networking investments. Each quarter, Packet delivers coverage of the latest industry trends, technology breakthroughs, and Cisco products and solutions, as well as network deployment and troubleshooting tips, configuration examples, customer case studies, certification and training information, and links to scores of in-depth online resources. You can access Packet magazine at this URL:

<http://www.cisco.com/packet>

- *iQ Magazine* is the quarterly publication from Cisco Systems designed to help growing companies learn how they can use technology to increase revenue, streamline their business, and expand services. The publication identifies the challenges facing these companies and the technologies to help solve them, using real-world case studies and business strategies to help readers make sound technology investment decisions. You can access iQ Magazine at this URL:

<http://www.cisco.com/go/iqmagazine>

- *Internet Protocol Journal* is a quarterly journal published by Cisco Systems for engineering professionals involved in designing, developing, and operating public and private internets and intranets. You can access the Internet Protocol Journal at this URL:

<http://www.cisco.com/ipj>

- World-class networking training is available from Cisco. You can view current offerings at this URL:

<http://www.cisco.com/en/US/learning/index.html>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 1**

### **Fabric Manager Applications**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



# Installation and Configuration

---

The Cisco Fabric Manager is a set of network management tools that supports Secure Simple Network Management Protocol version 3 (SNMPv3). It provides a graphical user interface (GUI) that displays real-time views of your network fabrics, and lets you manage the configuration of Cisco MDS 9000 Family devices and third-party switches.

This chapter contains the following sections:

- [About Cisco Fabric Manager, page 1-1](#)
- [Installing the Management Software, page 1-5](#)
- [Upgrading the Management Software, page 1-8](#)
- [Downgrading the Management Software, page 1-9](#)
- [Launching the Management Software, page 1-9](#)
- [Integrating Cisco Fabric Manager with Other Management Tools, page 1-11](#)
- [Running Fabric Manager Behind a Firewall, page 1-11](#)
- [Uninstalling the Management Software, page 1-12](#)

## About Cisco Fabric Manager

The Cisco Fabric Manager provides an alternative to the command-line interface (CLI) for most switch configuration commands. For information on using the CLI to configure a Cisco MDS 9000 Family switch, refer to the *Cisco 9000 Family Configuration Guide* or the *Cisco 9000 Family Command Reference*.

In addition to complete configuration and status monitoring capabilities for Cisco MDS 9000 switches, Fabric Manager provides powerful Fibre Channel troubleshooting tools. These in-depth health and configuration analysis capabilities leverage unique MDS 9000 switch capabilities: Fibre Channel Ping and Traceroute.

The Cisco Fabric Manager includes these management applications:

- Fabric Manager (client and server)
- Device Manager
- Performance Manager
- Fabric Manager Web Services

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Manager Server

The Fabric Manager server component must be started before running Fabric Manager. On a Windows PC, the Fabric Manager server is installed as a service. This service can then be administered using the Windows Services in the Control Panel. Fabric Manager server is responsible for discovery of the physical and logical fabric, and for listening for SNMP traps, syslog messages, and Performance Manager threshold events. See [Chapter 2, “Fabric Manager Server.”](#)

## Fabric Manager Client

The Fabric Manager client component displays a map of your network fabrics, including Cisco MDS 9000 Family switches, third-party switches, hosts, and storage devices. The Fabric Manager Client provides multiple menus for accessing the features of the Fabric Manager Server. See [Chapter 3, “Fabric Manager Client.”](#)

## Fabric Manager Server Proxy Services

The Fabric Manager Client and Device Manager use SNMP to communicate with the Fabric Manager Server. In typical configurations, the Fabric Manager Server may be installed behind a firewall. The SNMP proxy service available in Cisco MDS SAN-OS Release 2.1(1a) or later provides a TCP-based transport proxy for these SNMP requests. The SNMP proxy service allows you to block all UDP traffic at the firewall and configure Fabric Manager Client to communicate over a configured TCP port.

Fabric Manager uses the CLI for managing some features on the switches. These management tasks are used by Fabric Manager and do not use the proxy services. Your firewall must remain open for CLI access for the following:

- external and internal loopback test
- flash files
- create cli user
- security - iscsi users
- quiese pc
- show image version
- show tech
- switch resident reports (syslog, accounting)
- zone migration
- show cores

See the [“Running Fabric Manager Behind a Firewall”](#) section on page 1-11.

## Device Manager

The Device Manager presents two views of a single switch.

- Device View displays a graphic representation of the switch configuration and provides access to statistics and configuration information.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Summary View displays a summary of xE ports (Inter-Switch Links), Fx ports (fabric ports), and Nx ports (attached hosts and storage) on the switch, as well as Fibre Channel and IP neighbor devices. Summary or detailed statistics can be charted, printed, or saved to a file in tab-delimited format.

See [Chapter 4, “Device Manager.”](#)

## Performance Manager

Performance Manager provides detailed traffic analysis by capturing data with SNMP. This data is compiled into various graphs and charts that can be viewed with any web browser. See [Chapter 32, “Performance Monitoring.”](#)

## Fabric Manager Web Services

The Fabric Manager Web Services allows operators to monitor and obtain reports for MDS events, performance, and inventory from a remote location using a web browser. For information on installing and using Fabric Manager Web Services, see [Chapter 5, “Fabric Manager Web Services.”](#)

# Cisco MDS 9000 Switch Management

The Cisco MDS 9000 Family of switches can be accessed and configured in many different ways and supports standard management protocols. [Table 1-1](#) lists the management protocols that Fabric Manager supports to access, monitor, and configure the Cisco MDS 9000 Family of switches .

**Table 1-1**      **Supported Management Protocols**

Management Protocol	Purpose
Telnet/SSH	Provides remote access to the CLI for a Cisco MDS 9000 switch.
FTP/SFTP/TFTP, SCP	Copies configuration and software images between devices.
SNMPv1, v2c, and v3	<p>Includes over 80 distinct Management Information Bases (MIBs). Cisco MDS 9000 Family switches support SNMP version 1, 2, and 3 and RMON V1 and V2. RMON provides advanced alarm and event management, including setting thresholds and sending notifications based on changes in device or network behavior.</p> <p>By default, the Cisco Fabric Manager communicates with Cisco MDS 9000 Family switches using SNMPv3, which provides secure authentication using encrypted user names and passwords. SNMPv3 also provides the option to encrypt all management traffic.</p>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 1-1 Supported Management Protocols (continued)**

Management Protocol	Purpose
HTTP	Includes HTTP for web browsers to communicate with Fabric Manager Web Services and for the distribution and installation of the Cisco Fabric Manager software. It is not used for communication between the Cisco Fabric Manager server and Cisco MDS 9000 Family switches.
ANSI T11 FC-GS-3	Provides Fibre Channel-Generic Services (FC-GS-3) in the defining management servers in the Fabric Configuration Server (FCS). Fabric Manager uses the information provided by FCS on top of the information contained in the Name Server database and in the Fibre Channel Shortest Path First (FSPF) topology database to build a detailed topology view and collect information for all the devices building the fabric.
XML/CIM	Includes CIM server support for designing storage area network management applications to run on Cisco MDS SAN-OS.

## Storage Management Solutions Architecture

Management services required for the storage environment can be divided into five layers, with the bottom layer being closest to the physical storage network equipment, and the top layer managing the interface between applications and storage resources.

Of these five layers of storage network management, Cisco Fabric Manager provides tools for device (element) management and fabric management. In general, the Device Manager is most useful for device management (a single switch), while Fabric Manager is more efficient for performing fabric management operations involving multiple switches.

Tools for upper-layer management tasks can be provided by Cisco or by third-party storage and network management applications. The following summarizes the goals and function of each layer of storage network management:

- Device management provides tools to configure and manage a device within a system or a fabric. You use device management tools to perform tasks on one device at a time, such as initial device configuration, setting and monitoring thresholds, and managing device system images or firmware.
- Fabric management provides a view of an entire fabric and its devices. Fabric management applications provide fabric discovery, fabric monitoring, reporting, and fabric configuration.
- Resource management provides tools for managing resources such as fabric bandwidth, connected paths, disks, I/O operations per second (IOPS), CPU, and memory. You can use Fabric Manager to perform some of these tasks.
- Data management provides tools for ensuring the integrity, availability, and performance of data. Data management services include redundant array of independent disks (RAID) schemes, data replication practices, backup or recovery requirements, and data migration. Data management capabilities are provided by third-party tools.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Application management provides tools for managing the overall system consisting of devices, fabric, resources, and data from the application. Application management integrates all these components with the applications that use the storage network. Application management capabilities are provided by third-party tools.

## In-Band Management and Out-of-Band Management

Cisco Fabric Manager requires an out-of-band (Ethernet) connection to at least one Cisco MDS 9000 Family switch. You need either mgmt0 or IP over Fibre Channel (IPFC) to manage the fabric.

### mgmt0

The out-of-band management connection is a 10/100 Mbps Ethernet interface on the supervisor module, labeled mgmt0. The mgmt0 interface can be connected to a management network to access the switch through IP over Ethernet. You must connect to at least one Cisco MDS 9000 Family switch in the fabric through its Ethernet management port. You can then use this connection to manage the other switches using in-band (Fibre Channel) connectivity. Otherwise, you need to connect the mgmt0 port on each switch to your Ethernet network.

Each supervisor module has its own Ethernet connection; however, the two Ethernet connections in a redundant supervisor system operate in active or standby mode. The active supervisor module also hosts the active mgmt0 connection. When a failover event occurs to the standby supervisor module, the IP address and media access control (MAC) address of the active Ethernet connection are moved to the standby Ethernet connection.

### IPFC

You can also manage switches on a Fibre Channel network using an in-band IP connection. The Cisco MDS 9000 Family supports RFC 2625 IP over Fibre Channel, which defines an encapsulation method to transport IP over a Fibre Channel network.

IPFC encapsulates IP packets into Fibre Channel frames so that management information can cross the Fibre Channel network without requiring a dedicated Ethernet connection to each switch. This feature allows you to build a completely in-band management solution.

## Installing the Management Software

To install the software for the first time, or if you want to update or reinstall the software, access the supervisor module with a web browser. Click the **Install** links on the web page that is displayed. The software running on your workstation is verified to make sure you are running the most current version of the software. If it is not current, the most recent version is downloaded and installed on your workstation.



#### Note

Before upgrading or uninstalling Fabric Manager or Device Manager, make sure any instances of these applications have been shut down.

Installation options include:

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Upgrade/Downgrade - The installer detects your current version of Fabric Manager and Device Manager, and it provides the option to upgrade or downgrade. The default is to upgrade to the latest version of Fabric Manager or Device Manager.
- Uninstall - If you are downgrading from Cisco MDS SAN-OS 2.x to Cisco MDS SAN-OS 1.3x or earlier, use the Uninstall batch file or shell script. Do not delete the MDS 9000 folder as this might prevent your installation from being upgraded in the future.


**Note**

We recommend that you install the latest version of the Fabric Manager applications. Fabric Manager is backward-compatible with the Cisco MDS SAN-OS software running on the switches. When upgrading, upgrade the Fabric Manager software first, and then upgrade the Cisco MDS SAN-OS software on the switch.

## Before You Install

Before you can access the Cisco Fabric Manager, you must complete the following tasks:

- Install a supervisor module on each switch that you want to manage.
- Configure the supervisor module with the following values using the setup routine or the CLI:
  - IP address assigned to the mgmt0 interface
  - SNMP credentials (v1/v2 communities, or v3 user name and password), maintaining the same username and password for all the switches in the fabric

Cisco MDS SAN-OS Release 2.1(1a) or later supports AAA authentication using RADIUS, TACACS+ or local SNMP users.

The Cisco Fabric Manager software executable files reside on each supervisor module of each Cisco MDS 9000 Family switch in your network. The supervisor module provides an HTTP server that responds to browser requests and distributes the software to Windows or UNIX network management stations.

The Cisco Fabric Manager management software has been tested with the following software:

- Operating Systems
  - Windows 2000, 2003, XP
  - Solaris 2.8
  - Redhat Linux 7.2
- Java
  - Sun JRE and JDK 1.4.0, 1.4.1 and 1.4.2
  - Java Web Start 1.2 and 1.0.1
- Browsers
  - Internet Explorer 5.5 or later
  - Netscape 6 or later
  - Mozilla 1.0 or later



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Installation Procedure

To download and install the software on your workstation, follow these steps:

- Step 1** Enter the IP address or host name of the supervisor module in the Address or Location field of your browser.
- When you connect to the server for the first time, it checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. If not, a link is provided to the appropriate Sun Microsystems web page so you can install it. Fabric Manager looks for version 1.4(x) during installation.
- The supervisor module HTTP server displays the installation window.

- Step 2** Click the link to the Sun Java Virtual Machine software (if required) and install the software.
- Using the instructions provided by the Sun Microsystems website, reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.



**Note** Fabric Manager requires Java version 1.4(x). We recommend using Java version 1.4.2. To change the JRE version, start Java Web Start and set the Java preferences. In some cases, license validation from Cisco partners requires Java version 1.4.2\_04 or later. If you cannot install licenses from a Cisco partner, check to make sure your Java version is at least 1.4.2\_04.



**Note** You can run CiscoWorks on the same PC as Fabric Manager, even though the Java requirements are different. When installing the later Java version for Fabric Manager, make sure it does not overwrite the earlier Java version required for CiscoWorks. Both versions of Java can coexist on your PC.

- Step 3** Click on the desired installation link (**Fabric Manager**, **Device Manager**, or **Fabric Manager Web Services and Performance Manager**).
- Step 4** Select an installation folder for Fabric Manager on your workstation, as shown in [Figure 1-1](#). The default location is C:\Program Files\Cisco Systems\MDS 9000 for Windows. On a UNIX (Solaris or Linux) machine, the installation path name is /usr/local/cisco\_mds9000 or \$HOME/cisco\_mds9000, depending on the permissions of the user doing the installation.
- Step 5** Check the **Use Global Aliases in place of FC Aliases** if you want to use global device aliases or replace existing per VSAN FC aliases with global device aliases.

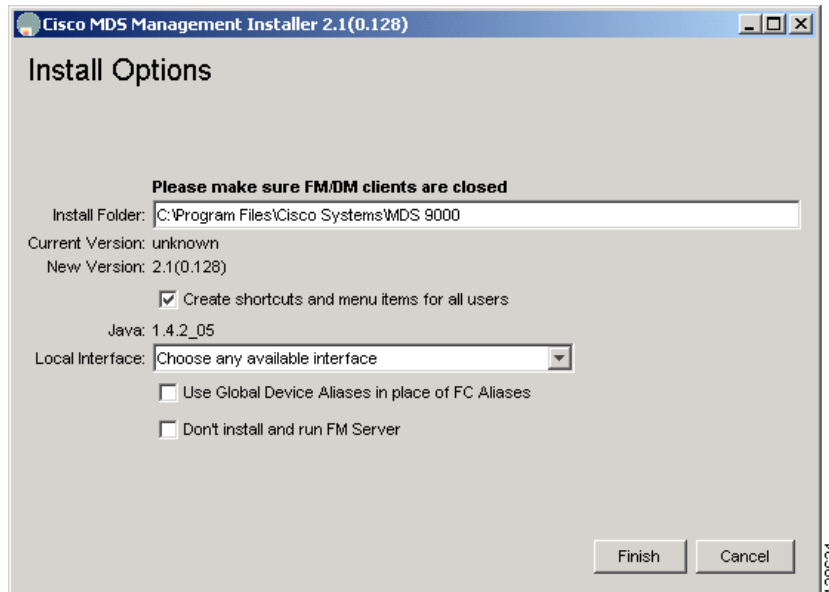


**Tip** After installation, you can choose to use global aliases by setting **fabric.globalAlias** to true in the server.properties file.

- Step 6** Check the **don't install and run FM Server** if you are installing just the Fabric Manager client on a remote workstation.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Figure 1-1** *Install Options*



A Cisco MDS 9000 program group is created under Start > Programs on Windows. This program group contains shortcuts to batch files in the install directory. Three services are started: Fabric Manager Server, Database, and Web Server. The Performance Manager server is installed but the service is not started at install time, as certain setup steps must be completed first.

On a UNIX (Solaris or Linux) machine, shell scripts are created in the install directory. The shell scripts that run the programs equivalent to the Windows services are: FMServer.sh, FMPersist.sh, PMCollector.sh and FMWebClient.sh. All server-side data and Performance Manager data are stored under the install directory.

Fabric Manager client cannot run without the server component, Fabric Manager Server. The server component is downloaded and installed when you download and install Fabric Manager or Device Manager. On a Windows machine you install the Fabric Manager Server as a service. This service can then be administered using Services in the Microsoft Windows Control Panel. The default setting for the Fabric Manager Server service is that the server is automatically started when the machine is rebooted. You can change this behavior by modifying the properties in Services.

## Upgrading the Management Software

If you log into a switch with Fabric Manager or Device Manager, and that switch has a later version of the management software, you are prompted to install the later version. To upgrade the Cisco MDS Fabric Manager software, follow the instructions described in the [“Installing the Management Software” section on page 1-5](#). You can also upgrade the software at any time by entering the IP address or host name of the supervisor module with the later version of software in the Address (Location) field of your browser.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Downgrading the Management Software

You can manage switches on your fabric with a later version of Fabric Manager than the version of Cisco MDS SAN-OS on the switches, but you cannot manage all features on switches with a later version of Cisco MDS SAN-OS.

### Downgrading to Cisco MDS SAN-OS Release 2.x or Later

To downgrade from any Cisco MDS SAN-OS Release 2.x or later to an earlier version of Cisco MDS SAN-OS Release 2.x, follow these steps:

- 
- Step 1** Close all instances of Fabric Manager Client or Device Manager on your workstation.
  - Step 2** Enter the IP address or host name of the supervisor module with the lower version of software in the Address or Location field of your browser and follow the installation steps. See the [“Installing the Management Software”](#) section on page 1-5.
- 

### Downgrading to Cisco MDS SAN-OS Release 1.3(x) or Earlier

To downgrade the management software from any Cisco MDS SAN-OS Release 2.x or later to Cisco MDS SAN-OS Release 1.3(x) or earlier, follow these steps:

- 
- Step 1** Close all instances of Fabric Manager Client or Device Manager on your workstation.
  - Step 2** Choose **Start > Programs > Cisco MDS 9000 > Uninstall** to uninstall Fabric Manager on Windows. Type `/usr/local/cisco_mds9000/uninstall.sh` or `$HOME/cisco_mds9000/uninstall.sh` to uninstall Fabric Manager on UNIX, depending on where Fabric Manager was installed.
  - Step 3** Enter the IP address or host name of the supervisor module with the lower version in the Address or Location field of your browser.
  - Step 4** Click on the desired installation link (**Fabric Manager**, **Device Manager**, or **Fabric Manager Web Services and Performance Manager**).
  - Step 5** Select an installation folder for Fabric Manager on your workstation.
- 

Unless you specify a different directory on a Windows PC, the version 1.3(x) software is installed in the default location of `.\Documents and Settings\USER_ID\cisco_mds9000`. A Cisco MDS program group is created under **Start > Programs**. On a UNIX (Solaris or Linux) machine, the installation path name is `/usr/local/.cisco_mds9000` or `$HOME/.cisco_mds9000`, depending on the permissions of the user doing the installation.

## Launching the Management Software

To launch the Fabric Manager (Fabric View) or Device Manager (Device View and Summary View), follow these steps:

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 1** Double-click the **Fabric Manager** icon or the **Device Manager** icon on your desktop or select the option from the Windows Start menu.
- If you started Fabric Manager, the Fabric Manager server loads. You see a log-in screen for Fabric Manager or Device Manager. (You briefly see a command-line window).
- Step 2** Click **Options** to expand the login screen if necessary to select the seed switch and SNMP configuration.
- Step 3** Enter the IP address or device name in the Device Name(s) field, or select an IP address from the list of previously accessed devices from the drop-down menu to the right of the Device Name(s) field.
- Step 4** Leave the **SNMPv3** check box checked to select SNMP version 3. Otherwise, uncheck the check box to use SNMP version 2.



**Note** The default authentication digest used for storing user names and passwords is MD5. In case you selected SHA instead, the relative check box in the Fabric Manager initial login screen should be checked.

- Step 5** Enter a **User Name** and **Password**.
- Step 6** If the SNMPv3 Privacy option is enabled, enter the **Privacy Password** used for encrypting management traffic.

The Privacy option causes all management traffic to be encrypted while, with SNMPv3, user names and passwords are always encrypted.



**Note** You can create users with management traffic encryption (so a Privacy password is required) or no management traffic encryption (no Privacy password is required). Requiring a Privacy password, and making it different from the authentication password, ensures stronger security.

- Step 7** Check the **Use SNMP Proxy** check box if you want Fabric Manager client to communicate with Fabric Manager Server through a TCP-based proxy server. See the [“Fabric Manager Server Proxy Services” section on page 1-2](#).



**Note** Load From Database check box should remain checked for normal operation. Uncheck this only if you require a fresh rediscovery of the fabric because of changed switch IP addresses. You may experience problems with out of sync SAN IDs in Fabric Manager if you uncheck this check box.

- Step 8** Click **Open**.
- You see either the Fabric Manager or the Device Manager.



**Note** When logging into Fabric Manager or Device Manager, the local SNMP database is checked first. If no username entry is found, the AAA database is checked.



**Note** If you have an inconsistent view of your fabric, rediscover the fabric with a user that has a network administrator or network operator role.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

# Integrating Cisco Fabric Manager with Other Management Tools

You can use Fabric Manager, Device Manager, and Performance Manager with other management tools. Here is a brief description of these tools. For more information on these tools and how they work together with the Cisco Fabric Manager management applications, see [Chapter 34, “Troubleshooting Your Fabric,”](#)

- Cisco Traffic Analyzer—Allows you to break down traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.
- Cisco Protocol Analyzer—Enables you to examine actual sequences of Fibre Channel frames easily using the Fibre Channel and SCSI decoders Cisco developed for Ethernet.
- Cisco Port Analyzer Adapter 2—Encapsulates SPAN traffic (both Fibre Channel control and data plane traffic) in an Ethernet header for transport to a Windows PC or workstation for analysis. Both the Cisco Traffic Analyzer and Cisco Protocol Analyzer required the PAA to transport MDS SPAN traffic to a Windows PC or workstation.

## Running Fabric Manager Behind a Firewall

For Windows PCs running Fabric Manager, Device Manager, and Performance Manager behind a firewall, certain ports need to be available.

The Fabric Manager client component and Device Manager use the first available UDP port for receiving SNMP responses. The UDP SNMP Trap local ports are (1162 for Fabric Manager, and 1163 or 1164 for Device Manager). Fabric Manager Client also opens TCP RMI port 9099. If Device Manager is opened from the Fabric Manager client, it listens on the first available UDP port for Fabric Manager requests.

The Fabric Manager Server proxy services feature, available in Cisco MDS SAN-OS Release 2.1(1a) or later, uses a configurable TCP port (9189 by default) for SNMP communications between the Fabric Manager client or Device Manager and Fabric Manager Server.

The Fabric Manager server component requires two predictable TCP ports to be opened on the firewall for an incoming connection:

- `java.rmi.registry.port = 9099`
- `java.rmi.server.remoteObjectPort = 9199`

As long as these two ports are opened, the Fabric Manager client can connect to the server. There may be other TCP ports connected to Fabric Manager Client, but they are initiated by the server, which is behind the firewall.

Below is a list of all ports used by the Fabric Manager applications:

Common to all applications

- SSH 22 (TCP)
- TELNET 23 (TCP)
- HTTP 80 (TCP)
- TFTP 69 (UDP)
- SYSLOG 514 (UDP)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Fabric Manager Server and Performance Manager

- SNMP\_TRAP 2162 (UDP)
- SNMP picks a random free local port (UDP) - (can be changed in server.properties)
- Java RMI 9099, 9199 to 9299 (TCP)

Fabric Manager Client

- Java RMI 9099, 9199 to 9299 (TCP)
- SNMP picks a random free local port. (UDP) or 9189 (TCP) if SNMP proxy is enabled (can be changed in server.properties)

Device Manager

- SNMP\_TRAP 1163 to 1170 (UDP) (picks one available in this range)
- SNMP picks a random free local port (UDP) or 9189 (TCP) if SNMP Proxy is enabled (can be changed in server.properties)

## Uninstalling the Management Software

To uninstall the Fabric Manager applications on a Windows PC, follow these steps:

**Step 1** Close all running instances of Fabric Manager and Device Manager.

**Step 2** Select **Start > Programs > Cisco MDS 9000 > Uninstall** to run the uninstall.bat script.

You can also run the batch file (located in the C:\Program Files\Cisco Systems\MDS 9000 folder by default) directly from the command line.



**Note** For older installations, delete the .cisco\_mds9000 folder. You will have to manually delete all desktop icons and program menu items.

On a Windows PC, this folder is created under the Documents and Settings folder (for example, d:\Documents and Settings\Administrator\.cisco\_mds9000 if you had installed it as user Administrator). On a UNIX machine, the default installation folder is /usr/bin.



**Note** You cannot downgrade from Fabric Manager Release 2.x to Fabric Manager Release 1.3(x). If you want to run Fabric Manager Release 1.3(x) on a PC that is running Fabric Manager Release 2.x, you must first uninstall Release 2.x and then install Release 1.3. Fabric Manager will not work if you have Release 2.x and Release 1.3 installed on the same PC.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To uninstall the Fabric Manager applications on a UNIX machine, follow these steps:

- 
- Step 1** For all releases starting with Release 2.x, run the shell script `$HOME/cisco_mds9000/Uninstall.sh` or `/usr/local/cisco_mds9000/uninstall.sh`, depending on where Fabric Manager was installed.
- Step 2** For all releases starting with Release 1.3(1), run the shell script `$HOME/.cisco_mds9000/Uninstall.sh` or `/usr/local/.cisco_mds9000/uninstall.sh`, depending on where Fabric Manager was installed.
- Step 3** For earlier installations, delete the `$HOME/.cisco_mds9000` folder.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Fabric Manager Server

---

Fabric Manager Server is a platform for advanced MDS monitoring, troubleshooting, and configuration capabilities. No additional software must be installed. The server capabilities are an integral part of the Cisco Fabric Manager software.

This chapter contains the following sections:

- [Fabric Manager Server Overview, page 2-1](#)
- [Fabric Manager Server Features, page 2-2](#)
- [Installing and Configuring Fabric Manager Server, page 2-2](#)
- [Fabric Manager Server Fabric Monitoring and Removal, page 2-7](#)
- [Fabric Manager Server Properties File, page 2-8](#)

## Fabric Manager Server Overview

Install Cisco Fabric Manager Server on a computer that you want to provide centralized MDS management services and performance monitoring. SNMP operations are used to efficiently collect fabric information, so Cisco Fabric Manager Server does not bog down your IP network with management traffic. The Cisco Fabric Manager software, including the server components requires about 20 MB of hard disk space on your workstation. Cisco Fabric Manager Server runs on Windows 2000, 2003, XP, Solaris 8.x or later, and Red Hat Linux.

Each computer configured as a Cisco Fabric Manager Server can monitor multiple Fibre Channel SAN fabrics. Up to 16 clients (by default) can connect to a single Cisco Fabric Manager Server concurrently. The Cisco Fabric Manager clients can also connect directly to an MDS switch in fabrics that are not monitored by a Cisco Fabric Manager Server, which ensures you can manage any of your MDS devices from a single console.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Manager Server Features

Cisco Fabric Manager Server has the following features:

- **Multiple Fabric Management**—Fabric Manager Server monitors multiple physical fabrics under the same user interface. This facilitates managing redundant fabrics. A licensed Fabric Manager Server maintains up-to-date discovery information on all configured fabrics so device status and interconnections are immediately available when you open the Fabric Manager client.

**Note**

The unlicensed Cisco Fabric Manager can only monitor and configure one fabric at a time. You must use the Open menu to switch to a new fabric, which causes the application to stop monitoring the previous one and to rediscover the new fabric.

- **Continuous Health Monitoring**—MDS health is monitored continuously, so any events that occurred since the last time you opened the Fabric Manager client are captured.
- **Roaming User Profiles**—The licensed Fabric Manager Server uses the roaming user profile feature to store your preferences and topology map layouts on the server, so that your user interface will be consistent regardless of what computer you use to manage your storage networks.

**Note**

You must have the same release of Fabric Manager Client and Fabric Manager Server.

## Installing and Configuring Fabric Manager Server

**Note**

Prior to running Fabric Manager Server, you should create a special Fabric Manager administrative user on each switch in the fabric or on a remote AAA server. Use this user to discover your fabric topology. See the [“Best Practices for Discovering a Fabric”](#) section on page 7-3.

To install Fabric Manager Server and set the initial configuration, follow these steps:

- Step 1** Install Fabric Manager and Fabric Manager server on your workstation. See the [“Installing Fabric Manager Server”](#) section on page 2-3.
- Step 2** Set the seed switch. See the [“Setting the Seed Switch”](#) section on page 2-4.
- Step 3** Optionally, create flows and collections for Performance Manager to monitor your fabric. See the [“Configuring Flows and Collections with Performance Manager”](#) section on page 2-4.
- Step 4** Set Fabric Manager Server to continuously monitor the fabric. See the [“Fabric Manager Server Fabric Monitoring and Removal”](#) section on page 2-7.
- Step 5** Repeat [Step 2](#) through [Step 4](#) for each fabric that you want to manage through Fabric Manager Server.
- Step 6** Install Web Services. See the [“Installing Fabric Manager Web Services”](#) section on page 2-6.
- Step 7** Verify Performance Manager is collecting data. See the [“Verifying Performance Manager Collections”](#) section on page 2-6.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Installing Fabric Manager Server

When you install Fabric Manager, the basic version of the Fabric Manager Server (unlicensed) is installed with it. After you click the Fabric Manager icon, a dialog box opens and you can enter the IP address of a computer running the Fabric Manager Server component. If you do not see the Fabric Manager Server IP address text box, click **Options** to expand the list of configuration options. If the server component is running on your local machine, leave **localhost** in that field. If you try to run Fabric Manager without specifying a valid server, you are prompted to start the Fabric Manager Server locally.

On a Windows PC, you install the Fabric Manager Server as a service. This service can then be administered using Services in the Administrative Tools. The default setting for the Fabric Manager Server service is that the server is automatically started when the Windows PC is rebooted. You can change this behavior by modifying the properties in Services.



### Note

If your computer has multiple interface cards (NICs), choose a local interface that can reach the Fibre Channel network on the clients and on the server.

## Unlicensed Versus Licensed Fabric Manager Server

When you install Fabric Manager, the basic unlicensed version of Fabric Manager Server is installed with it. To get the licensed features, such as Performance Manager, remote client support, and continuously monitored fabrics, you need to buy and install the Fabric Manager Server package.

However, trial versions of these licensed features are available. To enable the trial version of a feature, you run the feature as you would if you had purchased the license. You see a dialog box explaining that this is a demo version of the feature and that it is enabled for a limited time.

If you are evaluating one of these Fabric Manager Server features and want to stop the evaluation period for that feature, you can do that using Device Manager. See the [“Fabric Manager Server Licensing” section on page 9-12](#).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Setting the Seed Switch

When you run Fabric Manager, you must select the seed switch that Fabric Manager uses.

To set the seed switch, follow these steps:

- 
- Step 1** Double-click the Fabric Manager client icon on your workstation. You see the Fabric Manager dialog box. Click **Options** if necessary to expand the optional settings in this dialog box.
  - Step 2** Set **FM Server** to the IP address where you installed the Fabric Manager Server or set it to **localhost** if you installed Fabric Manager Server on your local workstation.
  - Step 3** Set the **Fabric Seed Switch** to the MDS 9000 family switch that you want Fabric Manager to use.
  - Step 4** Set the username and password as required to start Fabric Manager Client.
- 

## Configuring Flows and Collections with Performance Manager

If you are managing your fabrics with Performance Manager, you need to set up an initial set of flows and collections on the fabric. See the [“Historical Performance Monitoring” section on page 32-2](#) for a full description on Performance Manager features.

To create a flow in Performance Manager, follow these steps:

- 
- Step 1** Choose **Performance > Create Flows** to launch the wizard.
  - Step 2** Choose the **VSAN** from which you want to create flows. Flows are defined per VSAN.
  - Step 3** Click the **Type** radio button for the flow type you want to define.
  - Step 4** Check the **Clear old flows on modified switches** check box if you want to remove old flow data.
  - Step 5** Click **Next** to review the available flows for the chosen VSAN. Remove any flows you are not interested in.
  - Step 6** Click **Finish** to create the flow.
- 

## Using the Performance Manager Configuration Wizard

To create a collection using the Performance Manager Configuration Wizard in Fabric Manager, follow these steps:

- 
- Step 1** Choose **Performance > Create Collection** to launch the Performance Manager Configuration Wizard.
  - Step 2** Choose the VSANs from which you want to collect data or choose **All** to collect statistics across all VSANs in the fabric.
  - Step 3** Check the **Type** check boxes for each type of links flow or SAN element that you want included in your collection.
  - Step 4** If you want to ignore flows with zero counter values, check that check box.
  - Step 5** If you are using Cisco Traffic Analyzer, enter the URL where it is located on your network.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 6** Click **Next** to review the collection specification data. Remove any links, flows, or SAN elements you are not interested in.
  - Step 7** Click **Next** to configure other collection options.
  - Step 8** Check the appropriate check boxes if you want to include errors and discards in your collection, and if you want to interpolate data for missing statistics.
  - Step 9** Check the **Send event if traffic exceeds threshold** check box if you want to configure threshold events as explained in the [“Using Performance Thresholds” section on page 32-4](#).
  - Step 10** Click the **Use absolute values** radio button if you want absolute value thresholds or click the **Baseline values over** radio button if you want baseline thresholds.
  - Step 11** Choose the time window for baseline calculations if baseline thresholds are configured.
  - Step 12** Choose the Critical and Warning threshold values as a percent of link capacity (for absolute value thresholds ) or weighted average (for baseline thresholds).
  - Step 13** Click **Finish** to create the collection configuration file. You see a dialog box asking if you want to restart Performance Manager.
  - Step 14** Click **Yes** to restart Performance Manager to use this new configuration file, or click **No** to exit the Performance Manager Configuration Wizard without restarting Performance Manager. If you choose No, Performance Manager will not use the new configuration file until you restart it by choosing **Performance Manager > Collector > Restart**.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Installing Fabric Manager Web Services

You must install Fabric Manager Web Services to view Performance Manager reports through a web browser. See the [“Installing Fabric Manager Web Services” section on page 5-5](#).

To download and install the software on your workstation, follow these steps:

- 
- Step 1** Enter the IP address or host name of the supervisor module in the Location or Address field of your browser. You see the installation page displayed by the HTTP server of the supervisor module.
- When you connect to the server for the first time, it checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. If you do not have the correct version installed, a link is provided to the appropriate web page on the Sun Microsystems website so you can install it.
- Click the **Sun Java Virtual Machine** software link (if required) to install the software.
  - Using the instructions provided by the Sun Microsystems website, reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.

**Note**

Fabric Manager requires Java version 1.4(x). We recommend Java version 1.4.2. To change the Java Runtime Environment (JRE) version, start Java Web Start and set the Java preferences.

---

- Step 2** Click the **Fabric Manager Web Services** installation link. You see a prompt asking for permission to install the application on your workstation.
- Step 3** Click **Yes** to run the installer, which detects the installed version of the software, and prompts for upgrades/downgrades and other options if applicable.

**Note**

If TCP port 80 is in use, Fabric Manager Web Services checks port 8080 next. If that port is also in use, Fabric Manager Web Services uses the next available port. You can set the TCP port that you want Fabric Manager Web Services to use during the installation process.

---

## Verifying Performance Manager Collections

Once Performance Manager collections have been running for five or more minutes, you can verify that the collections are gathering data by choosing **Performance Manager > Reports** in Fabric Manager. You see the first few data points gathered in the graphs and tables.

**Note**

Viewing reports requires installing Fabric Manager Web Services. See the [“Installing Fabric Manager Web Services” section on page 5-5](#).

---

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Fabric Manager Server Fabric Monitoring and Removal

Fabric Manager Server can continuously monitor a fabric, whether or not an instance of Fabric Manager (client) is monitoring that fabric. A continuously monitored fabric is automatically reloaded and monitored by Fabric Manager Server after the server starts up. Fabrics that are monitored by Fabric Manager Server can have their data managed by Performance Manager. Both the Continuous Monitor feature and Performance Manager require the Fabric Manager Server license. However, you can “check out” these features without a license for a limited time.

### Designating a Fabric for Continuous Monitoring

When you quit the Fabric Manager client, you are prompted as to whether or not you would like to have Fabric Manager Server continuously monitor that fabric. Alternatively, you can use Fabric Manager client to select a fabric to monitor.

To continuously monitor a fabric, follow these steps:

- 
- Step 1** From Fabric Manager, select **Server > Admin**. You see a list of fabrics in the Server Admin dialog box.
  - Step 2** Check the **Continuously Monitor** check box next to the fabric(s) you want Fabric Manager Server to monitor.
  - Step 3** Click **Apply**.

The Continuously Monitor feature requires the purchase of the Fabric Manager Server license package. If you have not purchased and installed this package, you see a popup window informing you that you are about to enable a demo license for this feature. Click **Yes** to enable the demo license.



**Note** When you are finished trying the licensed features, you can “check in” the feature by clicking the **Check In FM** button as described in the [“Fabric Manager Server Licensing”](#) section on [page 9-12](#).

---

- Step 4** Click **Close** to close the Server Admin dialog box.
- 



**Note** If you will be collecting data on these fabrics using Performance Manager, you should now configure flows and define the data collections. These procedures are described in [Chapter 31, “Network Monitoring.”](#)

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Removing a Fabric from Monitoring

To remove a fabric from the Fabric Manager Server monitoring list, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | From Fabric Manager, choose <b>Server &gt; Admin</b> . You see a list of fabrics in the Server Admin dialog box.           |
| <b>Step 2</b> | Uncheck the <b>Continuously Monitor</b> check box next to the fabrics you no longer want Fabric Manager Server to monitor. |
| <b>Step 3</b> | Click <b>Apply</b> .   |
| <b>Step 4</b> | Click <b>Close</b> to close the Server Admin dialog box.   |
- 

## Fabric Manager Server Properties File

The Fabric Manager Server properties file (MDS 9000\server.properties) contains a list of properties that determine how the Fabric Manager Server will function. You can edit this file with a text editor, or you can set the properties through the Fabric Manager Web Services GUI, under the Admin tab.

The server properties file contains these five general sections:

- **RMI SPECIFIC**—Contains the settings for the RMI registry.
- **SNMP SPECIFIC**—Contains the settings for SNMP requests, responses, and traps.
- **SNMP PROXY SERVER SPECIFIC**—Contains the settings for SNMP proxy server configuration and TCP port designation.
- **GLOBAL FABRIC**—Contains the settings for fabrics, such as discovery and loading.
- **CLIENT SESSION**—Contains the settings for Fabric Manager clients that can log into the server.
- **EVENTS**—Contains the settings for syslog messages.

The following are new or changed server properties for Cisco MDS SAN-OS Release 2.x:

- **fabric.globalAlias**—Specifies whether Fabric Manager Server should discover aliases from a global alias server (deviceAlias) or a VSAN-based alias server (fcAlias). Global aliases of a fabric are fetched if this value is set to **true** and a manageable global alias server exists in the fabric.
- **java.rmi.data.portRange**—Specifies the TCP port range that RMI uses to send/receive data. The starting port number is also the Fabric Manager Server port, so that it should always be set to a non-zero value. The port range is also used in configuring TCP port access on a firewall. This property replaces the java.rmi.server.remoteObjPort property in earlier releases.
- **fabric.autoReload**—Specifies whether to automatically reload persistent fabrics from DB when server starts up. The default is true if unspecified.
- **fabric.loadFromDB**—Specifies whether to load fabric from DB when it is opened. The default is false if unspecified.
- **proxy.autostart**—Specifies whether to automatically start SNMP proxy server when Fabric Manager Server starts. Default is true if unspecified. Note, proxy will not be started if snmp.proxy is specified as a non-localhost.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- proxy.localaddress—Specifies the local network interface name, e.g. "eth0", "eth1", or a local IP address, to bind proxy server socket on a multi-homed localhost. If IP address 0.0.0.0 is specified, proxy will be bound to the wildcard address, an IP address chosen by the kernel.
- proxy.localport—Specifies the local TCP port to listen for client connection. If port 0 is specified, proxy will be bound to a port chosen by the kernel. Default is 9198.

For more information on setting the server properties, see [Chapter 5, “Fabric Manager Web Services.”](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Fabric Manager Client

---

The Cisco Fabric Manager Client is a java based GUI application that provides easy access to the Fabric Manager applications from a remote workstation.

This chapter contains the following sections:

- [Fabric Manager Client Overview, page 3-1](#)
- [Launching Fabric Manager Client, page 3-2](#)
- [Using Fabric Manager Client, page 3-3](#)
- [Setting Fabric Manager Preferences, page 3-12](#)
- [Network Fabric Discovery, page 3-14](#)
- [Modifying Device Grouping, page 3-14](#)
- [Control of Administrator Access with Users and Roles, page 3-15](#)
- [Fabric Manager Wizards, page 3-15](#)
- [Fabric Manager Troubleshooting Tools, page 3-16](#)

## Fabric Manager Client Overview

The Cisco Fabric Manager is a Java and SNMP-based network fabric and device management tool with a GUI that displays real-time views of your network fabric, including Cisco MDS 9000 and third-party switches, hosts, and storage devices.

In addition to complete configuration and status monitoring capabilities for Cisco MDS 9000 switches, Fabric Manager client provides powerful Fibre Channel troubleshooting tools. These in-depth health and configuration analysis tools leverage unique MDS 9000 switch capabilities including Fibre Channel ping and traceroute.



### Note

---

You must have the same release of Fabric Manager Client and Fabric Manager Server.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Manager Advanced Mode

Cisco MDS SAN-OS Release 2.1(1a) introduces Advanced Mode. Advanced Mode is enabled by default and provides the full suite of Fabric Manager features, including security, IVR, iSCSI, and FICON. Uncheck the **Advanced** check box in the upper right corner of Fabric Manager client to simplify the user interface. In this mode, you can access the basic MDS 9000 features like VSANs, zoning, and configuring interfaces.

## Launching Fabric Manager Client

To launch Fabric Manager client, follow these steps:

- 
- Step 1** Double-click the **Fabric Manager** icon and follow the instructions described in [“Launching the Management Software”](#) section on page 1-9 to launch Fabric Manager Client from your desktop,.
- Step 2** To launch Fabric Manager Client from within a running instance of Fabric Manager, follow these steps:
- Choose **Open Fabric** from the Fabric Manager **File** menu.
  - Click the **Open Switch Fabric** button from the Fabric Manager toolbar.
  - Select the IP address of the fabric seed switch you want to access. If you do not see the fabric seed switch in the pop-up window, click the **Options** button to expand the pop-up window options.
  - Click **Open** if the fabric you want to open has the same username and password as the fabric you already have open.

Fabric Manager displays the new fabric and adds a tab to the Fabric pane.

- Step 3** Click each fabric’s tab to view the fabric.

If the fabric you want to open has a different username and password, enter the username and password and click **Open**. You are prompted for whether or not you want the open fabric(s) to be monitored in the background. That fabric is then closed on the Fabric Manager client, and the new fabric is opened.




---

**Note** Changes made using Fabric Manager are applied to the running configuration of the switches you are managing. If you have made changes to the configuration or performed an operation (such as activating zones), Fabric Manager prompts you to save your changes before you exit.


---

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Using Fabric Manager Client

This section describes the Fabric Manager client interface, as shown in [Figure 3-1](#).

**Figure 3-1 Fabric Manager Main Window**



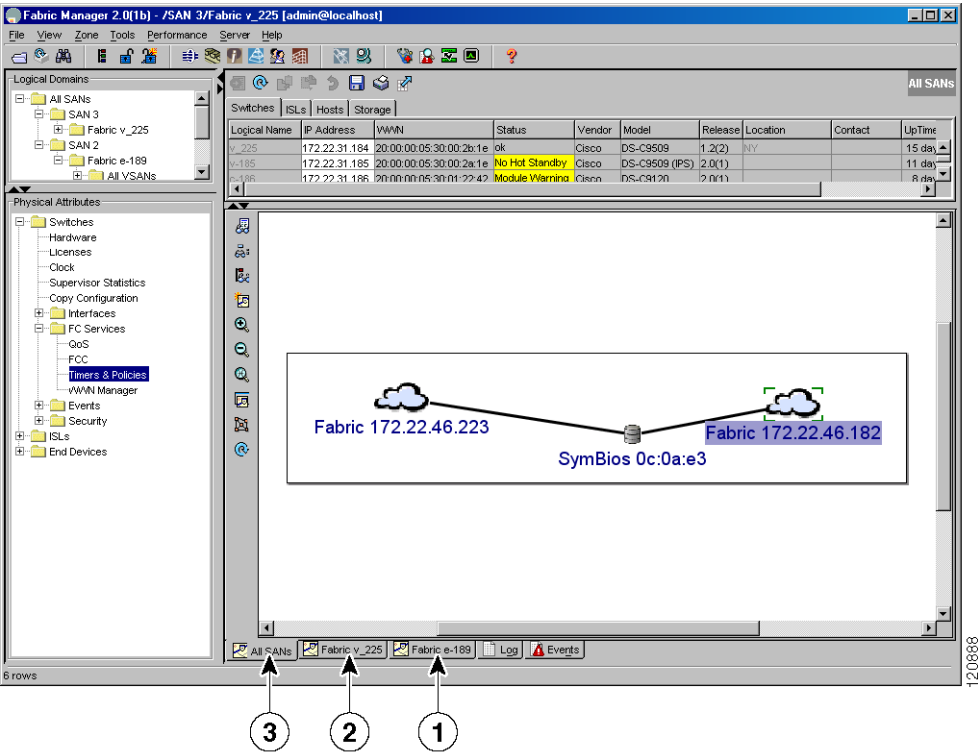
1	Menu bar—Provides access to options that are organized by menus.
2	Toolbar—Provides icons that provide direct access to the most commonly used options on the File, Tools, and Help menus.
3	Information pane—Displays information about whatever option is selected in the menu tree.
4	Status Bar (right side)—Shows the last entry displayed by the discovery process, and the possible error message.
5	Fabric pane—Displays a map of the network fabric, including switches, hosts, and storage. It also provides tabs for displaying log and event data.
6	Logical domains—Displays a tree of configured SANs, fabrics, VSANs and zone.
7	Physical attributes—Displays a tree of available configuration tasks depending on the SAN, fabric, VSAN, or zone selected above. Lists the switches and end devices in the logical selection.
8	Status Bar (left side)—Shows short-term transient messages, such as the number of rows displayed in a table.

Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).

# Multiple Fabric Display

Fabric Manager can display multiple fabrics in the same pane (see [Figure 3-2](#)).

Figure 3-2 Fabric Manager’s Multiple Fabric Display



1	The Fabric view tab for fabric v_255. When selected, the Fabric view displays fabric v_255.
2	The Fabric view tab for fabric e_189. When selected, the Fabric view displays fabric e_189.
3	All SANS tab (selected), showing two fabrics.

**Note**

The same username and password must be used to log into multiple fabrics.

The information for both fabrics is displayed, with no need to select a seed switch. To see details of a fabric, select the tab for that fabric at the bottom of the Fabric pane, or double-click on the cloud icon for the fabric in the All SANS tab.

## Contents Panes













The following sections describe the panes in the Fabric Manager view. You can resize each pane by dragging the boundaries between each region or by clicking the **Minimize** or **Maximize** controls.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Pane







The Fabric pane shows the graphical representation of your fabric. [Table 3-1](#) explains the graphics you may see displayed, depending on which devices you have in your fabric.

**Table 3-1**      ***Fabric Manager Graphics***

Icon or Graphic	Description
	Director class MDS 9000.
	Non-director class MDS 9000.
	Generic Fibre Channel switch.
	Cisco SN5428.
	An orange line through a device indicates that the device is manageable but there are operational problems.
	An orange "X" through a device or link indicates that the device or ISL is not working properly.
	A red line through a device indicates that the device is not manageable.
	A red "X" through a device or link indicates that the device is down or that the ISL is down.
	Fibre Channel HBA (or enclosure).
	Fibre Channel target (or enclosure).
	iSCSI host.
	Fibre Channel ISL and edge connection.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 3-1**      **Fabric Manager Graphics (continued)**

Icon or Graphic	Description
	Fibre Channel PortChannel.
	IP ISL and edge connection.
	IP PortChannel.
	Fibre Channel loop (storage).
	IP cloud (hosts). This icon is also used to represent a fabric when viewing a SAN (multiple fabrics) in the Fabric Manager Fabric pane.
	Any device, cloud, or loop with a box around it means that there are hidden links attached.

If a switch or director is grayed out, Fabric Manager can no longer communicate with it.

There are multiple tabs on the bottom of the Fabric pane:

- Fabric—When displaying multiple fabrics, each fabric has its own tab. You can switch between fabrics by clicking on their respective tabs.
- Log—Displays messages that describe Fabric Manager operations, such as fabric discovery.
- Events—Displays information about the SNMP traps received by the management station. This includes combination events as detected by discovery and important traps like license, SNMP, and FICON.

When viewing large fabrics in the Fabric pane, it is helpful to:

- Turn off end device labels
- Collapse loops
- Collapse expanded multiple links (collapsed multiple links are shown as very thick single lines)
- Dim or hide portions of your fabric by VSAN



**Note**

When a VSAN, zone, or zone member is selected in the VSAN tree, the map highlighting changes to identify the selected objects. To remove this highlighting, click the **Clear Highlight** button on the Fabric pane toolbar or choose **Clear Highlight** from the pop-up menu.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Main Menu







The menu bar at the top of the Fabric Manager main window provides options for managing and troubleshooting the current fabric and for controlling the display of information on the Fabric pane. The menu bar provides the following menus:

- **File**—Opens a new fabric, rediscovers the current fabric, locates switches, sets preferences, prints the map, and clears (right-click on log) or exports the Fabric pane log.
- **View**—Changes the appearance of the map (these options are duplicated on the Fabric pane toolbar).
- **Zone**—Manages zones, zone sets, and inter-VSAN routing (IVR).
- **Tools**—Verifies and troubleshoots connectivity and configuration, as described in the “[Fabric Manager Troubleshooting Tools](#)” section on page 3-16.
- **Performance**—Runs and configures Performance Manager and Cisco Traffic Analyzer, and generate reports.
- **Server**—Runs administrative tasks on clients and fabrics. Provides Fabric Manager Server management and a purge command. Lists switches being managed.
- **Help**—Displays online help topics for specific dialog boxes in the Information pane.

## Toolbar

The Fabric Manager main toolbar provides buttons for accessing the most commonly used menu bar options as shown in [Table 3-2](#).

**Table 3-2**      ***Fabric Manager Client Main Toolbar***

Icon	Description
	Open switch fabric.
	Rediscover current fabric.
	Find in the map.
	Create VSAN.
	Launch DPVM wizard.
	Edit full zone database.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 3-2**      ***Fabric Manager Client Main Toolbar (continued)***











Icon	Description
	Launch IVR zone wizard.
	Launch PortChannel wizard.
	Launch FCIP wizard.
	Launch iSCSI wizard.
	Launch QoS wizard.
	Configure users and roles.
	Launch IP-ACL wizard.
	Launch License Install wizard.
	Launch Software Install wizard.
	Perform switch health analysis.
	Perform fabric configuration analysis.
	Perform end-to-end connectivity analysis.
	Monitor ISL performance.
	Show on-line help.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Information Pane

The Information pane displays tables of information associated with the option selected from the menu tree in the Logical Domains or Physical Attributes panes. The Information pane toolbar provides buttons for performing one or more of the operations shown in [Table 3-3](#).

**Table 3-3 Information Pane Toolbar**

Icon		Description
	Apply Changes	Applies configuration changes.
	Refresh Values	Refreshes table values.
	Create Row	Opens the appropriate dialog box to create a new row in the table.
	Delete Row	Deletes the currently highlighted rows from the table.
	Copy/Ctrl+C	Copies data from one row to another.
	Paste/Ctrl +V	Pastes the data from one row to another.
	Undo Changes/Ctrl-Z	Undoes the most recent change.
	Export	Exports and saves information to a file.
	Print Table	Prints the contents of the Information pane.
	Detach Table	Displays a non-editable copy of the table in the Information pane in its own window, which you can move around the screen.



### Note

After making changes you must save the configuration or the changes will be lost when the device is restarted.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



**Note**

The buttons that appear on the toolbar vary according to the option you select. They are activated or deactivated (dimmed) according to the field or other object that you select in the Information pane.

## Logical Domains Pane

Use the Logical Domains pane to manage attributes for SANs, fabrics, VSANs, and zones.

To manage these things, right-click one of the folders in the tree and click a menu item from the pop-up menu. You see the appropriate configuration dialog box.

The default name for the fabric is the name, IP address, or WWN for the principal switch in VSAN 1. If VSAN 1 is segmented, the default name is chosen from a principal switch with the smallest WWN. In order, the fabric names you may see are:

- Fabric <sysName>
- Fabric <ipAddress>
- Fabric <sWWN>

## Physical Attributes Pane

Use the Physical Attributes pane to display a tree of the options available for managing the switches in the currently discovered SAN, fabric, VSAN, or zone.

To select an option, click a folder to display the options available and then click the option. You see the table with information for the selected option in the Information pane. The Physical Attributes pane provides the following main folders:

- Switches—Views and configures hardware, system, licensing, and configuration files.
- Interfaces—Views and configures FC Physical, FC Logical, Ethernet, SVC, and PortChannel interfaces.
- FC Services—Views and configures Fibre Channel network configurations.
- IP—Views and configures IP storage and IP services.
- Events—Views and configures events, alarms, thresholds, notifications, and informs.
- Security—Views and configures MDS management and FC-SP security.
- ISLs—Views and configures Inter-Switch Links.
- End Devices—Views and configures end devices.

## Status Bar

The status bar at the bottom of the Fabric Manager window shows the last entry displayed by the discovery process, and the possible error message on the right side. The status bar displays a message stating that something has changed in the fabric and a new discovery is needed. The status bar shows both short-term, transient messages (such as the number of rows displayed in the table), and long-term discovery issues.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Context Menus

When you right-click an icon in the Fabric pane, you see a pop-up menu with options that vary depending on the type of icon selected. The various options available for different objects include the following:

- Open an instance of Device Manager for the selected switch.
- Open a CLI session for the selected switch.
- Copy the display name of the selected object.
- Execute a **ping** or **traceroute** command for the device.
- Show or hide end devices.
- View attributes
- Quiesce and disable members for PortChannels
- Set the trunking mode for an ISL.
- Create or add to a PortChannel for selected ISLs.

The Fabric pane has its own toolbar with options for saving, printing, and changing the appearance of the map. When you right-click on the map, a pop-up menu appears that provides options (duplicated on the toolbar) for changing the appearance of the map.



### Note

You can launch web-based or non-web-based applications from the Fabric pane. To do this, you assign an IP address to the storage port or enclosure. Then right-click to bring up the pop-up menu, and select **Device Manager**.

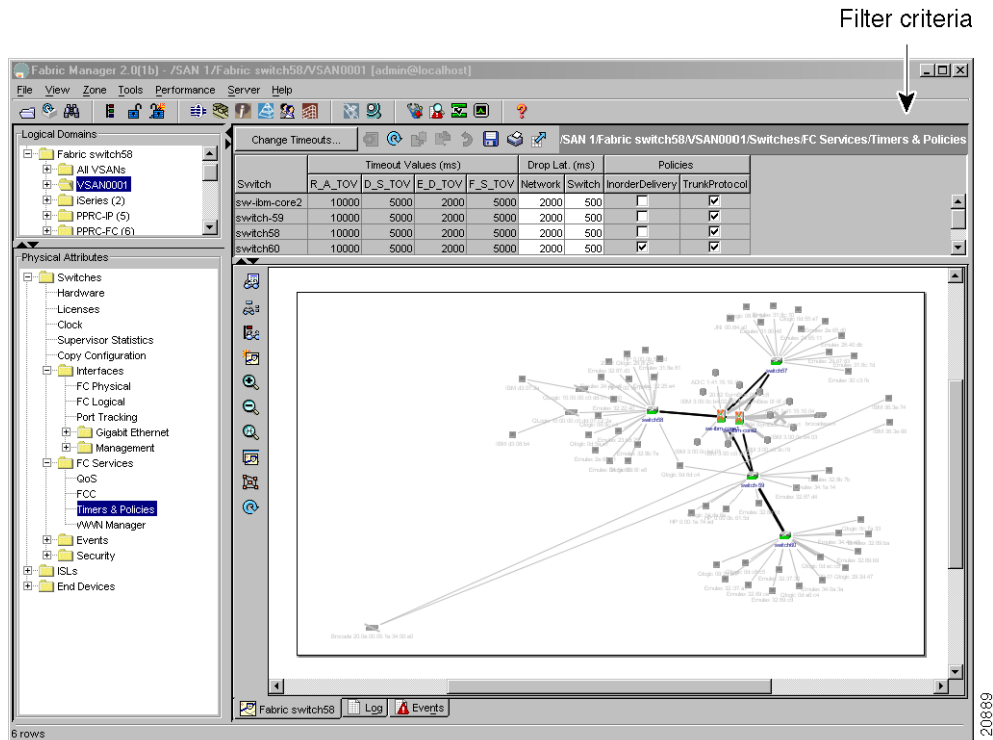
## Filtering

Fabric Manager has a built-in filtering mechanism that displays only the data you are interested in. To filter, first select the SAN, fabric, and VSAN from the Logical Domains pane. This narrows the scope of what is displayed in the Fabric pane. Any information that does not belong to the selected items is dimmed. Also, any information that does not belong to the selected items is not displayed in the tables in the Information pane. As shown in [Figure 3-3](#), the filter you select is displayed at the top right of the Fabric Manager window.

To further narrow the scope, select attributes from the Physical Attributes pane. The Fabric Manager tables, display, and filter criteria change accordingly.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 3-3 Fabric Manager's Filtering Mechanism**



## Detachable Tables

Cisco MDS SAN-OS Release 2.0(2b) introduced detachable tables in Fabric Manager. You can, for example, detach tables and move them to different areas on your desktop so you can compare similar tables from different VSANs. Or, you can keep informational tables open from one view while you examine a different area in Fabric Manager. To detach tables, click the **Detach Table** icon in the Information pane in Fabric Manager.

## Setting Fabric Manager Preferences

To set your preferences for the behavior of the Fabric Manager, choose **File > Preferences** from the Fabric Manager menu bar. You see the Preferences dialog box with the following tabs for setting different components of the application:

- General
- SNMP
- Map

The default General preferences for Fabric Manager are:

- Show Switch Name by—Displays the switches in the Fabric pane by IP address, DNS name, or logical name. The default setting for this value is Logical Name.

## ***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **Show WorldWideName (WWN) Vendor**—This displays the world wide name vendor name in any tables or listings displayed by Fabric Manager. If **Prepend Name** is checked, the name is displayed in front of the IP Address of the switch. If **Replacing Vendor Bytes** is checked, the name is displayed instead of the IP address. The default setting is enabled (checked) with the **Prepend Name** option.
- **Show End Device Using**—Displays end devices in the Fabric pane using alias or pWWN alias. The default setting for this value is Alias.
- **Append Enclosures to End Device Names**—The default setting for this value is OFF.
- **Show Shortened iSCSI Names**—The default setting for this value is OFF.
- **Show Timestamps as Date/Time**—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- **Telnet Path**—The path for the telnet.exe file on your system. The default is telnet.exe, but you will need to browse for the correct location.



### **Note**

If you browse for a path or enter a path and you have a space in the pathname (for example, c:\program files\telnet.exe), then the path will not work. To get the path to work, you must manually place quotes around it (for example, "c:\program files\telnet.exe").

- **Use Secure Shell instead of Telnet**—Specifies whether to use SSH or Telnet when using the CLI to communicate with the switch. If enabled, you must specify the path to your SSH application. The default setting is disabled.
- **Confirm Deletion**—Displays a confirmation pop-up when you delete part of your configuration using Fabric Manager. The default setting is enabled (checked).
- **Export Tables with Format**—Specifies the type of file that is created when you export a table using Device Manager. The options are tab-delimited or XML. The default setting is Tab-Delimited.

The default SNMP preferences for Fabric Manager are:

- **Retry request 1 time(s) after 5 sec timeout**—You can set the retry value to 0-5, and the timeout value to 3-30.
- **Trace SNMP packets in Log**—The default setting for this value is OFF.
- **Enable Audible Alert when Event Received**—The default setting for this value is OFF.

The default Map preferences for Fabric Manager are:

- **Display Unselected VSAN Members**—Displays the unselected VSAN members in the Fabric pane. The default setting for this value is ON.
- **Display End Devices**—Displays the fabric's end devices in the Fabric pane. The default setting for this value is ON.
- **Display End Device Labels**—Displays the fabric's end device labels in the Fabric pane. The default setting for this value is ON.
- **Expand Loops**—Displays the loops in the fabric as individual connections in the Fabric pane. The default setting for this value is OFF.
- **Expand Multiple Links**—Displays multiple links in the Fabric pane as separate lines rather than as one thick line. The default setting for this value is ON.
- **Open New Device Manager Each Time**—Opens a new instance of Device Manager each time you invoke it from a switch in your fabric. The default value is OFF, which means only one instance of Device Manager is open at a time.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **Select Switch or Link from Table**—Allows you to select a switch or link in the Fabric pane by clicking on the switch or link in a table in the information pane. The default setting for this value is disabled (unchecked), which means clicking on a switch or link in the table does not change the switch or link selection in the Fabric pane.
- **Layout New Devices Automatically**—Automatically places new devices in the Fabric pane in an optimal configuration. The default setting for this value is OFF. In this mode, when you add a new device, you must manually reposition it if the initial position does not suit your needs.
- **Use Quick Layout when Switch has  $\geq 30$  End Devices**—The default setting for this value is 30. You can enter any number in this field. Enter **0** to disable Quick Layout.
- **Override Preferences for Non-default Layout**—The default setting for this value is ON.
- **Automatically Save Layout**—If this option is enabled, any changes in layout are automatically saved. The default setting for this value is ON.
- **Detach Overview Window**—Allows you to more easily center the Fabric pane on the area of the fabric you want to see. (This is most useful for large fabrics that cannot be displayed entirely within the Fabric pane.) Bring up the overview window by clicking the **Show/Hide Overview Window** button. It overlays the fabric window and remains there until you click the **Show/Hide Overview Window** button again. If you enable this preference, you can detach the overview window and move it to one side while you access the Fabric pane. The default setting for this value is disabled (unchecked).

## Network Fabric Discovery

Cisco Fabric Manager collects information on the fabric topology through SNMP queries to the switches connected to Fabric Manager. The switch replies after having discovered all devices connected to the fabric by using the information coming from its FSPF technology database and the Name Server database, and collected using the Fabric Configuration Server's request/response mechanisms defined by the FC-GS-3/4 standard. When you start the Fabric Manager, you enter the IP address (or host name) of a "seed" switch for discovery.

After you start Fabric Manager and discovery completes, Fabric Manager presents you with a view of your network fabric, including all discovered switches, hosts, and storage devices.

## Modifying Device Grouping

Because not all devices are capable of responding to FC-GS-3 requests, different ports of a single server or storage subsystem may be displayed as individual end devices on the Fabric Manager map.

To group end devices in a single enclosure to have them represented by a single icon on the map, follow these steps:

- 
- Step 1** Select **Storage** or **Hosts** from the Fabric Manager's Physical tree in the Navigation pane.  
You see the end devices displayed in the Information pane.
  - Step 2** Click one of the devices or the Name field that you want to be in the enclosure.
  - Step 3** Enter a name to identify the new enclosure's icon on the Fabric Manager Fabric pane.
  - Step 4** Click once on the Name field for that device. To select more than one name, hold down the **Shift** key and click each of the other names.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 5** Press **Ctrl-C** to copy the selected name(s).

**Step 6** Press **Ctrl-V** to paste the name into the Name field for that device.



**Note**

To remove devices from an enclosure, triple click on the name of the device and press **Delete**. To remove an enclosure, repeat this step for each device in the enclosure.

## Control of Administrator Access with Users and Roles

Cisco MDS 9000 Family switches support role-based management access whether using the CLI or the Cisco Fabric Manager. This lets you assign specific management privileges to particular roles and then assign one or more users to each role.

Cisco Fabric Manager uses SNMPv3 to establish role-based management access. After completing the setup routine, a single role, user name, and password are established. The role assigned to this user allows the highest level of privileges, which includes creating users and roles. Use the Cisco Fabric Manager to create roles and users, and to assign passwords as required for secure management access in your network.

## Fabric Manager Wizards

Fabric Manager client provides a series of wizards to facilitate common configuration tasks. These wizards include:

- VSAN—Creates VSANs on multiple switches in the fabric and sets VSAN attributes including interop mode, load balancing, and FICON.
- Zone Edit Tool —Creates zone sets, zones, and aliases. Adds members to zones, and edits zone database.
- IVR Zone—Creates IVR zone sets, zones, and aliases. Enables IVR NAT and auto-topology. Adds members to IVR zones, and edits IVR zone database.
- PortChannel—Creates PortChannels from selected ISLs either manually or automatically. Sets PortChannel attributes like channel ID and trunking mode.
- FCIP —Creates FCIP links between Gigabit Ethernet ports. Enables Fibre Channel Write Acceleration and IP compression
- DPVM—Establishes dynamic port VSAN membership, enables auto learning, and activates the DPVM database.
- iSCSI—Zones iSCSI initiators and adds VSAN to target allowed VSAN list.
- QoS—Sets QoS attributes for zones in the selected VSAN.
- IP ACL—Creates ordered IP access control lists and distributes to selected switches in the fabric.
- License Install—Facilitates download and installation of licenses in selected switches in the fabric.
- Software Install—Verifies image compatibility and installs software images on selected switches in the fabric.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Manager Troubleshooting Tools

Fabric Manager has several troubleshooting tools available from the toolbar or Tools menu. Procedures for using these tools are described in [Chapter 34, “Troubleshooting Your Fabric.”](#) Here is a brief description of each tool.

- **Zone Merge Analysis**—The zone merge analysis tool (available from the Zone menu) enables you to determine if zones will merge successfully when two Cisco MDS switches are interconnected. If the interconnected switch ports allow VSANs with identical names or contain zones with identical names, then Fabric Manager verifies that the zones contain identical members. The merge analysis tool can be run before attempting a merge or after fabrics are interconnected to determine zone merge failure causes.
- **End-to-End Connectivity**—Fabric Manager’s end-to-end connectivity analysis tool uses FC Ping to verify interconnections between Cisco MDS switches and end-device (HBAs and storage devices) in a particular VSAN. In addition to basic connectivity, Fabric Manager can optionally verify that:
  - Latency is less than a specified value.
  - Paths are redundant.
  - Zones contain at least two members.

End devices are connected to a manageable switch (have an currently active in-band or out-of-band management path.)

- **Switch Health Analysis**—You can run an in-depth switch health analysis with Fabric Manager. It verifies the status of all critical Cisco MDS switches, modules, ports, and Fibre Channel services. Over 40 conditions are checked per switch port. This tool provides a very fast, simple, and thorough way to assess Cisco MDS switch health.
- **Fabric Configuration Analysis**—Fabric Manager includes a fabric configuration analysis tool. It compares the configurations of all Cisco MDS switches in a fabric to a reference switch or a policy file. You can define what functions to check and what type of checks to perform. The analysis can look for mismatched values, and missing or extra values. If all configuration checking is performed for all functions, over 200 checks are performed for each Cisco MDS switch.

After the analysis is run, the results are displayed with details about the issues that were discovered. You can automatically resolve configuration differences by selecting them and clicking the **Resolve** button. Fabric Manager automatically changes the configuration to match the reference switch or policy file.



## Device Manager

---

This chapter contains descriptions of, and instructions for using, the Cisco MDS 9000 Device Manager.

This chapter contains the following sections:

- [Device Manager Overview, page 4-1](#)
- [Device Manager Features, page 4-1](#)
- [Launching Device Manager, page 4-2](#)
- [Using Device Manager, page 4-3](#)
- [Setting Device Manager Preferences, page 4-8](#)

## Device Manager Overview

Device Manager provides a graphic representation of a Cisco MDS 9000 Family switch chassis, including the installed switching modules, the supervisor modules, the status of each port within each module, the power supplies, and the fan assemblies.

The tables in the Fabric Manager Information pane basically correspond to the dialog boxes that appear in Device Manager. However, while Fabric Manager tables show values for one or more switches, a Device Manager dialog box shows values for a single switch. Also, Device Manager provides more detailed information for verifying or troubleshooting device-specific configuration than what is available from Fabric Manager.

## Device Manager Features

Device Manager provides two views: Device View and Summary View. You can use Summary View to monitor all of the ports on the switch. You can use the Device View to perform any switch-level configuration task including the following:

- Manage ports, PortChannels, and trunking.
- Manage SNMPv3 security access to switches.
- Manage CLI security access to switch.
- Manage alarms, events, and notifications.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Save and copy configuration files and software image.
- View hardware configuration.
- View chassis, module, port status and statistics.

## Launching Device Manager

You can launch Device Manager in two ways.

To launch Device Manager from your desktop, double-click the **Device Manager** icon and follow the instructions described in the “[Launching the Management Software](#)” section on page 1-9.

To launch Device Manager from Fabric Manager, follow these steps:


- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Right-click the switch you want to manage on the Fabric Manager Fabric pane and click <b>Device Manager</b> from the pop-up menu that appears. |
| <b>Step 2</b> | Double-click a switch in the Fabric Manager Fabric pane.   |
| <b>Step 3</b> | Select a switch in the Fabric Manager Map pane and choose <b>Tools &gt; Device Manager</b> .   |
-

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Using Device Manager

This section describes the Device Manager interface, as shown in [Figure 4-1](#).

**Figure 4-1** Device Manager, Device Tab



1	Menu bar	5	Status
2	Toolbar	6	Supervisor modules
3	Tabs	7	Switching or services modules
4	Legend		

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Menu Bar




The menu bar at the top of the Device Manager main window provides options for managing and troubleshooting a single switch. The menu bar provides the following options:

- **Device**—Opens an instance of Device Manager, sets management preferences, sets the page layout, opens a Telnet/SSH session with the current switch, and closes the Device Manager application.
- **Physical**—Allows you to view and manage inventory, modules, temperature sensors, power supplies, fans, and the entire system.
- **Interface**—Allows you to configure and manage PortChannels, as well as Fibre Channel, Ethernet, iSCSI, and FICON ports. Also provides diagnostic, management and monitoring capabilities, as well as SPAN and port tracking.
- **FC**—Allows you to configure and manage VSAN, domain, and name server characteristics. Also provides advanced configuration capabilities.
- **FICON**—Allows you to configure and manage FICON VSANs, configure RLIR ERL information, and swap selected FICON ports.
- **IP**—Allows you to configure and manage the following types of information: FCIP, iSCSI, iSNS, routes, VRRP, and CDP.
- **Security**—Allows you to configure and manage FC SP, port security, iSCSI security, SNMP security, common roles, SSH, AAA, and IP ACLs.
- **Admin**—Allows you to save, copy, edit, and erase the switch configuration, monitor events, manipulate Flash files, manage licenses, configure NTP, use CFS, and reset the switch. Also enables you to use the **show tech support**, **show cores**, and **show image** commands.
- **Logs**—Shows the various logs: message, hardware, events, and accounting. Also displays FICON link incidents, and allows you to configure the syslog setup.
- **Help**—Displays online help topics for specific dialog boxes in the Information pane.

## Toolbar Icons










The Device Manager toolbar provides quick access to many Device Manager features. Once the icon is selected, a dialog box may open that allows configuration of the feature. The toolbar provides the main Device and Summary View icons as shown in [Table 4-1](#).

**Table 4-1**      **Device Manager Main Toolbar**

Icon	Description
 Open Device	Opens the Device Manager view for another switch, with the option to open this view in a separate window.
 Refresh Display	Communicates with the switch and displays the information in the Device Manager view.
 Command Line Interface	Opens a separate CLI command window to the switch.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 4-1**      **Device Manager Main Toolbar (continued)**

Icon		Description
	Configure Selected	Opens a configuration dialog box for the selected component (line card or port).
	SysLog	Opens a window that lists the latest system messages that occurred on the switch.
	Threshold Manager	Opens the Threshold Manager dialog box that provides statistical monitoring and event reporting for the switch.
	VSANs	Opens the VSAN dialog box that provides VSAN configuration for the switch.
	SNMP Users and Roles	Opens the SNMP configuration dialog box for SNMP users and roles.
	Save Configuration	Saves the current running configuration to the startup configuration.
	Copy	Copies configuration file between server and switch
	Toggle FICON/Interface Port Labels	Toggles the FICON and interface port labels.
	Help	Accesses online help for Device Manager.

## Dialog Boxes

If a toolbar icon is selected, a dialog box may open that allows configuration of the selected feature. The dialog box may include table manipulation icons. See the [“Information Pane” section on page 3-9](#) for descriptions of these icons.

## Tabs

Click the **Device** tab on the Device Manager main window to see a graphical representation of the switch chassis and components.

Click the **Summary** tab on the Device Manager main window to see a summary of active interfaces on a single switch, as well as Fibre Channel and IP neighbor devices. The Summary View also displays port speed, link utilization, and other traffic statistics. There are two buttons in the upper left corner of the Summary View tab used to monitor traffic. To monitor traffic for selected objects, click the **Monitor**

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Selected Interface Traffic Util%** button. To display detailed statistics for selected objects, click the **Monitor Selected Interface Traffic Details** button. You can set the poll interval, the type or Rx/Tx display, and the thresholds.

## Legend

The legend at the bottom right of the Device Manager indicates port status, as follows:

### Colors

- Green—The port is up.
- Brown—The port is administratively down.
- Red—The port is down or has failed.
- Amber—The port has a minor fault condition.
- Gray—The port is unreachable.

### Labels

- X—Link Failure
- E—ISL
- TE—Multi-VSAN ISL
- F—Host/Storage
- FL—F Loop
- I— iSCSI



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- SD—Span Destination
- CH—Channel
- CU—Control Unit

**Note**

For a detailed table describing the legend, see the [“There is a red/orange/dotted line through the switch. What’s wrong?”](#) section on page 35-14.

## Supervisor and Switching Modules

In the Device View, you can right-click on an object and get information on it, or configure it. If you right-click on a module, the menu shows the module number and gives you the option to configure or reset the module. If you right-click on a port, the menu shows the port number and gives you the option to configure, monitor, enable ,disable, set beacon mode, or perform diagnostics on the port.

**Tip**

You can select multiple ports in Device Manager and apply options to all the selected ports at one time. Either select the ports by clicking the mouse and dragging it around them, or hold down the **Control** key and click on each port.

To enable or disable a port, right-click the port and click **Enable** or **Disable** from the pop-up menu. To enable or disable multiple ports, drag the mouse to select the ports and then right-click the selected ports. Then click **Enable** or **Disable** from the pop-up menu.

To manage trunking on one or more ports, right-click the ports and click **Configure**. In the dialog box that appears, right-click the current value in the Trunk column and click **nonTrunk**, **trunk**, or **auto** from the pull-down list.

To create PortChannels using Device Manager, click **PortChannels** from the Interface menu. For detailed instructions, see [Chapter 17, “PortChannel Configuration.”](#) You can also use Fabric Manager to conveniently create a PortChannel.

**Note**

To create a PortChannel, all the ports on both ends of the link must have the same port speed, trunking type, and administrative state.

## Context Menus

Context menus are available in both Device Manager views by right-clicking on a device or table.

From Device View:

- Device—Right-click a device to bring up a menu that shows you what you clicked (system, module, or power supply) and to give you the option to configure or reset the device.
- Port— Right-click a port to bring up a menu that shows you the number of the port you have clicked, and to give you the option to configure, monitor, enable, disable, set beacon mode, or perform diagnostics on the port.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

From Summary View:

- Table—Right-click the table to show a list of which columns to display in that table: Interface, Description, VSANs, Mode, Connected To, Speed (Gb), Rx, Tx, Errors, Discards, and Log. Click the Description field to bring up the appropriate configuration dialog box for the port type.

## Setting Device Manager Preferences

To set your preferences for the behavior of the Device Manager application, choose **Device > Preferences** from the Device menu. You can set the following preferences:

- Retry requests x time(s) after x sec timeout—Allows you to set the retry request values. The default settings are 1 time after a 5-second timeout.
- Enable status polling every x secs—Allows you to set the status polling value. The default setting is enabled (checked) with a time of 40 seconds.
- Trace SNMP packets in Message Log—Allows you to set whether Device Manager traces SNMP packets and logs the trace. The default setting is disabled (unchecked).
- Register for Events after Open, listen on Port 1163—Allows you to register this switch so that events are logged once you open Device Manager. The default setting is enabled (checked).
- Confirm Deletion—Displays a popup confirmation when you delete part of your configuration using Device Manager. The default setting is enabled (checked).
- Show WorldWideName (WWN) Vendor—Displays the world wide name vendor name in any table or listing displayed by Device Manager. If **Prepend** is checked, the name is displayed in front of the IP address of the switch. If **Replace** is checked, the name is displayed instead of the IP address. The default setting is enabled (checked) with the **Prepend** option.
- Show Timestamps as Date/Time—Displays timestamps in the date/time format. If this preference is not checked, timestamps are displayed as elapsed time. The default setting is enabled (checked).
- Telnet Path—Sets the path for the telnet.exe file on your system. The default is telnet.exe, but you need to browse for the correct location.

**Note**

If you browse for a path or enter a path and you have a space in the pathname (for example, c:\program files\telnet.exe, then the path will not work. To get the path to work, manually place quotes around it (for example, "c:\program files\telnet.exe").

- Use Secure Shell instead of Telnet—Specifies whether to use SSH or Telnet when using the CLI to communicate with the switch. If enabled, you must specify the path to your SSH application. The default setting is disabled.
- CLI Session Timeout x secs (0= disable)—Specifies the timeout interval for a CLI session. Enter 0 to disable (no timeout value). The default setting is 30 seconds.
- Show Tooltips in Physical View—Determines whether tooltips are displayed in Physical (Device) View. The default setting is enabled (checked).
- Label Physical View Ports With:—Specifies the type of label to assign to the ports when you are in Physical (Device) View. The options are FICON and Interface. The default setting is Interface.
- Export Table—Specifies the type of file that is created when you export a table using Device Manager. The options are Tab-Delimited or XML. The default setting is Tab-Delimited.



## Fabric Manager Web Services

---

With Fabric Manager Web Services you can monitor Cisco MDS switch events, performance, and inventory from a remote location using a web browser. This chapter contains the following sections:

- [Fabric Manager Web Services Overview, page 5-1](#)
- [Installing Fabric Manager Web Services, page 5-5](#)
- [Launching and Using Fabric Manager Web Services, page 5-7](#)

### Fabric Manager Web Services Overview

Using Fabric Manager Web Services, you can monitor MDS switch events, performance, and inventory, and perform minor administrative tasks.

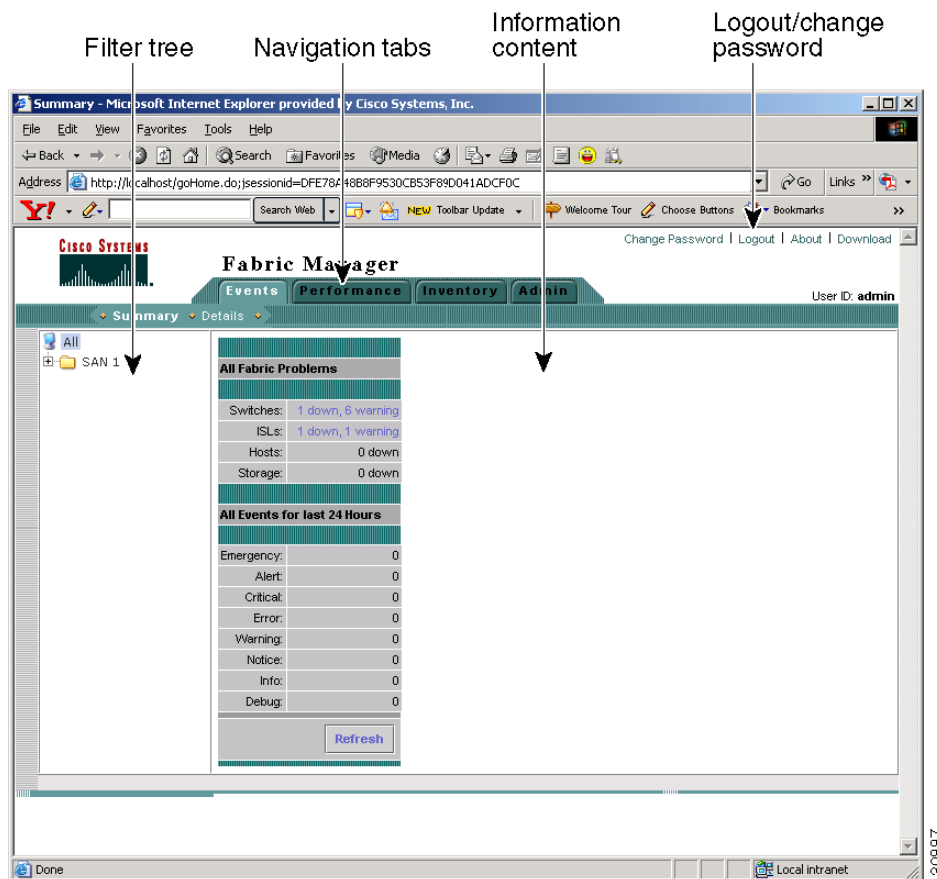
Fabric Manager Web Services provides the following features:

- **Summary and drill down reports**—The Performance Manager summary report provides a high-level view of your network performance. This report lists the average and peak throughput and provides hot-links to additional performance graphs and tables with additional statistics. Both tabular and graphical reports are available for all interconnections monitored by Performance Manager. Performance Manager also analyzes daily, weekly, monthly and yearly trends. These reports are only available if you create a collection using Performance Manager and start the collector. See the [“Historical Performance Monitoring” section on page 32-2](#).
- **Zero maintenance for statistics storage**—No maintenance is required to maintain Performance Manager’s round-robin database, because its size does not grow over time. At prescribed intervals the oldest samples are averaged (rolled-up) and saved. A full two days of raw samples are saved for maximum resolution. Gradually the resolution is reduced as groups of the oldest samples are rolled up together.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Fabric Manager Web Services displays in a web browser as shown in [Figure 5-1](#).

**Figure 5-1** *Fabric Manager Web Services.*



This section contains the following features:

- [Filter Tree, page 5-2](#)
- [Performance, page 5-4](#)
- [Inventory, page 5-4](#)
- [Admin, page 5-4](#)

## Filter Tree

Fabric Manager Web Services uses a filter navigation tree on the left pane to control the scope of the features in Fabric Manager Web Services. The filter tree expands or collapses based on clicking the + or - icons. Select the scope you want to access by expanding or collapsing the filter tree and then clicking on the file or folder that represents your desired scope. You can select All, or a specific SAN, fabric or VSAN from the filter tree in the left pane. The features accessible from the tabs is limited to the scope of what you select in the filter tree.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Events

The Events tab shows events and issues for the selected items, persistent across user sessions.

The events tab contains the following subtabs:

- Summary—Shows the summary of events and problems for All SANs, or a selected SAN, fabric, or switch. You can click on any of the blue links for more information about that item.
- Details—Shows a detailed list of events and hardware, accounting, or link incident switch syslogs. You can filter these events by severity, date, and type of event.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Performance

The Performance tab shows an overview of the average throughput and link utilization of SAN components. You see pie charts for the throughput and utilization. You can click on a pie chart to view a table of the data. In these tables, clicking on a blue link will display a graph of that data, if applicable. The Filter drop-down menu at the top right of the screen allows you to filter the data based on various periods of time.

The Performance tab contains the following subtabs:

- Summary—Shows the total utilization and throughput in summary form.
- Hosts—Shows a detailed list of host port traffic and errors.
- Storage - Shows a detailed list of storage port traffic and errors.
- ISLs—Shows a detailed list of ISL traffic and errors.
- Flows —Shows a detailed list of host-to-storage traffic.



### Note

Performance Manager shows statistics for fabrics that you have configured collections for using the Collection Wizard. See the [“Historical Performance Monitoring”](#) section on page 32-2.

## Inventory

The Inventory tab shows an inventory of the selected SAN, fabric, or switch. You can export this information to an ASCII file in comma-separated value format, that can be read by applications such as Microsoft Excel. You can set the number of rows and columns per page.

The Inventory tab contains the following subtabs:

- Summary—Shows VSANs, switches, ISLs, and ports.
- VSANs—Shows details about VSANs.
- Switches—Shows attributes (such as IP address, vendor, and model).
- Modules—Shows details for MDS switching and services modules, fans, and power supplies.
- ISLs—Shows the inter-switch links.
- Hosts—Shows the host ports.
- Storage—Shows the storage ports.
- Zones—Shows the active zone members (including those in inter-VSAN zones).

## Admin

The Admin tab allows you to perform minor administrative and configuration tasks on the Fabric Manager Server sending data to your web client.

The Admin tab contains the following subtabs:

- Status—Displays the status of, and allows you to start and stop the Database Server, Fabric Manager Server, and Performance Collector services on your server. You should only need to restart services if something is not working properly, or if too large a percentage of system resources are being consumed.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **Configure**—Allows you to configure various parameters for Fabric Manager Server.
- **Logs**—Allows you to view all the logs from the various services running on the Fabric Manager Server.
- **Web Users**—Allows you to configure Fabric Manager Web Services local or RADIUS user authentication.
- **Events**—Allows you to view the configuration settings for traps and syslog messages.

**Note**

If you see a database file lock error in the database log, you can fix it by shutting down and restarting the database server using the Web Client.

## Installing Fabric Manager Web Services

If you are installing the Fabric Manager Web Services software for the first time, or if you want to update or reinstall the software, you access the supervisor module of the switch using a web browser. Install Fabric Manager Web Services on the same workstation where you installed Fabric Manager Server.

To download and install the software on your workstation, follow these steps:

- Step 1** Enter the IP address or host name of the supervisor module in the Location or Address field of your browser. You see the installation page displayed by the HTTP server of the supervisor module.
- When you connect to the server for the first time, it checks to see if you have the correct Sun Java Virtual Machine version installed on your workstation. If you do not have the correct version installed, a link is provided to the appropriate web page on the Sun Microsystems website so you can install it.
- a. Click the **Sun Java Virtual Machine** software link (if required) to install the software.
  - b. Using the instructions provided by the Sun Microsystems website, reconnect to the supervisor module by reentering the IP address or host name in the Location or Address field of your browser.

**Note**

Fabric Manager requires Java version 1.4(x). We recommend Java version 1.4.2. To change the Java Runtime Environment (JRE) version, start Java Web Start and set the Java preferences.

- Step 2** Click the **Fabric Manager Web Services** installation link. You see a prompt asking for permission to install the application on your workstation.
- Step 3** Click **Yes** to run the installer, which detects the installed version of the software, and prompts for upgrades/downgrades and other options if applicable.

**Note**

If TCP port 80 is in use, Fabric Manager Web Services checks port 8080 next. If that port is also in use, Fabric Manager Web Services uses the next available port. You can set the TCP port that you want Fabric Manager Web Services to use during the installation process.

Unless you specify a different directory on a Windows PC, the software is installed in the default location of **C:\Program Files\Cisco Systems\MDS 9000**. A **Cisco MDS 9000** program group is created under **Start > Programs**. This program group contains shortcuts to Fabric Manager and Device manager.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

On a UNIX (Solaris or Linux) machine, the installation path is /usr/local/cisco\_mds9000. If this directory is not writable by the user, which is the case for non-root users, the default is set to \$HOME/cisco\_mds9000. Shell scripts are created in the bin directory.



#### Note

On a Windows PC, you install Fabric Manager Web Services as a service. This service can then be administered using the Services Panel from the Windows Control Panel. By default the Fabric Manager Web Services automatically starts when the workstation is rebooted. You can change this behavior by modifying the properties in the Services Panel.

## Using Fabric Manager Web Services with SSL

Fabric Manager Web Services uses TCP port 80 by default. If you want to install SSL certificates and use Fabric Manager Web Services over HTTPS (using TCP port 443 or another custom port), you need a certificate for each external IP address that accepts secure connections. You can purchase these certificates from a well-known Certificate Authority (CA).

To modify Fabric Manager Web Services to use SSL, follow these steps:

**Step 1** Stop Fabric Manager Web Services if you have already launched it. If you installed this on Windows, you can stop the service using Windows Services under Administrative Tools.

**Step 2** Open \tomcat\conf\server.xml from the directory that you installed Fabric Manager Web Services, using a text editor. You see the following lines in the beginning after some copyright information.:

```
<Connector className="org.apache.catalina.connector.http.HttpConnector"
    port="80" minProcessors="5" maxProcessors="75"
    enableLookups="false" redirectPort="8443"
    acceptCount="10" debug="0" connectionTimeout="60000"/>
<!-- Define an SSL HTTP/1.1 Connector on port 8443 -->
<!--
<Connector className="org.apache.catalina.connector.http.HttpConnector"
    port="8443" minProcessors="5" maxProcessors="75"
    enableLookups="true"
    acceptCount="10" debug="0" scheme="https" secure="true">
    <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
        clientAuth="false" protocol="TLS"/>
</Connector>
-->
```

**Step 3** Comment the first <Connector> element and uncomment the second one. Note that changes in port from 8443 to 443 and the addition of keystore and keypass. Your file should look like the following example:

```
<!--
<Connector className="org.apache.catalina.connector.http.HttpConnector"
    port="80" minProcessors="5" maxProcessors="75"
    enableLookups="false" redirectPort="8443"
    acceptCount="10" debug="0" connectionTimeout="60000"/>
-->
<!-- Define an SSL HTTP/1.1 Connector on port 8443 -->
<Connector className="org.apache.catalina.connector.http.HttpConnector"
    port="443" minProcessors="5" maxProcessors="75"
    enableLookups="true"
    acceptCount="10" debug="0" scheme="https" secure="true">
    <Factory className="org.apache.catalina.net.SSLServerSocketFactory"
        clientAuth="false" protocol="TLS"
        keystoreFile="C:\Program Files\Cisco Systems\MDS 9000\keystore"
        keystorePass="changeit"/>
```



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

</Connector>

- Step 4** Save this file.
- Step 5** Restart Fabric Manager Web Services.
- 

## Launching and Using Fabric Manager Web Services

Before you can use Fabric Manager Web Services to monitor a switch, the service must be started on the server you are connecting through. The browser does not have to be on the same workstation where Fabric Manager Web Services is installed.

To launch Fabric Manager Web Services, follow these steps:

- Step 1** If you are on the same workstation where you installed Fabric Manager Web Services, then open your browser and in the Location field enter **`http://localhost:PORT`**, where *PORT* is 80 by default. Enter your port number if you specified a different port during installation.

If you are on a different workstation from where you installed Fabric Manager Web Services, then open your browser and in the location field enter **`http://<yourServerAddress>:PORT`**, where *<yourServerAddress>* is the address where you installed Fabric Manager Web Services, and *PORT* is 80 by default. Enter your port number if you specified a different port during installation.



**Tip** Select the Windows **Start > Control Panel > Services** to verify that Fabric Manager Web Services is started or start Fabric Manager Web Services.  
On a UNIX workstation, use the following command:  
**`$ /usr/local/cisco_mds9000/bin/FMWebClient.sh status`**

---

You see the login screen for Fabric Manager Web Services (see [Figure 5-2](#)). The text field at the bottom shows the Message of the Day from the server you log into.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Figure 5-2** Fabric Manager Web Services Login Screen



**Step 2** Enter your user name and password.

**Step 3** Click the **Login** button.



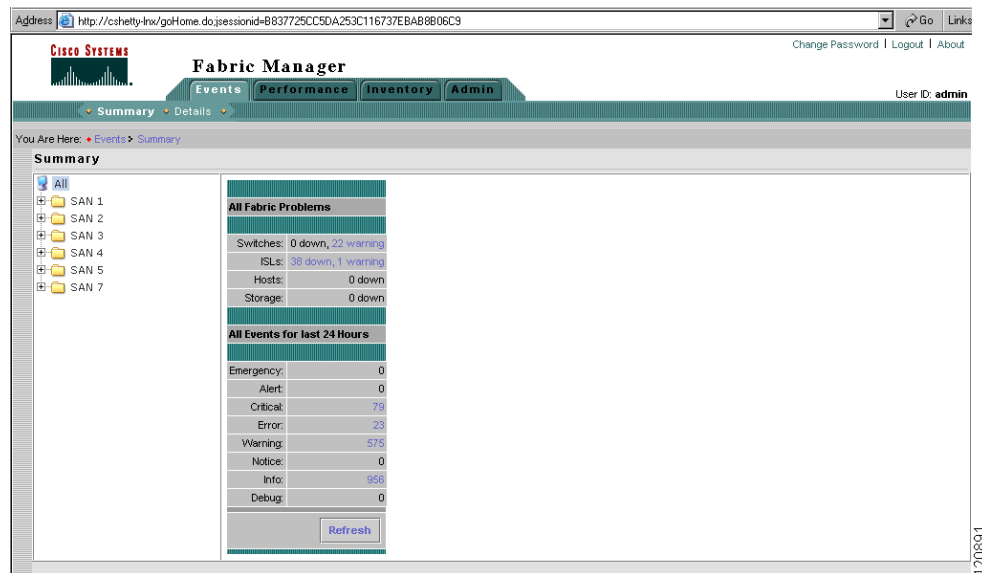
**Note**

When attempting to log in, you may see “No valid ID - server cannot find ID.” This happens if no discovery is done for that fabric using FM Server, so no username or password is recorded. To resolve this issue, open Fabric Manager and discover the fabric. Performance Manager will record the username and password. Then log in to Fabric Manager Web Services again.

After launching Fabric Manager Web Services, you see the initial screen, which is the Events > Summary screen (see [Figure 5-3](#)). Fabric Manager Web Services polls the Fabric Manager Server database to display the managed devices in the left pane.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Figure 5-3** *Events > Summary Screen*



## Monitoring Fabrics from Fabric Manager Web Services

Fabric Manager Web Services reports information gathered by Fabric Manager server on any fabrics known to Fabric Manager server.

To start or stop monitoring a fabric from Fabric Manager server using Fabric Manager Web Services, follow these steps:

- Step 1** Choose **Admin > Configure**. You see the Configuration options
- Step 2** Click **Fabrics** from the left navigation pane. You see the list of fabrics monitored by Fabric Manager server.
- Step 3** Select a fabric and click **Stop Monitoring** to discontinue data collection for that fabric.
- Step 4** Click Monitor... to start monitoring another fabric in Fabric Manager Web Services. You see the new fabric dialog box.
- Step 5** Set the seed switch, username and password options for this fabric and click **Monitor** to begin monitoring this fabric.

## Setting Up a Guest User

You may want a separate guest user account to access the reports available in Fabric Manager Web Services. You can create a case-insensitive guest user who can only view reports. The guest user cannot change the guest password, nor can the guest user access the Admin tab on Fabric Manager Web Services. You can set up a guest user either remotely in the AAA server, or locally using local authentication.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To create a guest user in the local database, follow these steps:

- 
- Step 1** Launch Fabric Manager Web Services. See the “[Launching and Using Fabric Manager Web Services](#)” section on page 5-7.
- Step 2** Choose **Admin > Web Users** to update the authentication used by Fabric Manager Web Services.
- Step 3** Click **Local Database**. You see the list of users in the local database.
- Step 4** Click **Add** to add a new user.
- Step 5** Set the username to **guest** and set the password.



---

**Note** The username guest is a reserved name (case insensitive). Fabric Manager Web services prevents the guest user from changing the guest password.

---

- Step 6** Click **Add** to add the guest user.
-



## Performance Manager

---

The primary purpose of Fabric Manager is to manage the network. A key management capability is network performance monitoring.

This chapter contains the following section:

- [Performance Manager Architecture, page 6-1](#)

## Performance Manager Architecture

Performance Manager gathers network device statistics historically and provides this information graphically using a web browser. It presents recent statistics in detail and older statistics in summary. Performance Manager also integrates with external tools such as Cisco Traffic Analyzer.

The Performance Manager has three operational stages:

- **Definition**—Uses two configuration wizards. The Flow Wizard sets up flows in the switches while the Collection Wizard create a collection configuration file.
- **Collection**—Reads the configuration file and collects the desired information.
- **Presentation**—Generates web pages to present the collected data.

Performance Manager can collect statistics for ISLs, hosts, storage elements, and configured flows. Flows are defined based on a host-to-storage (or storage-to-host) link. Performance Manager gathers statistics from across the fabric based on collection configuration files. These files determine which SAN elements and SAN links Performance Manager gathers statistics for. Based on this configuration, Performance Manager communicates with the appropriate devices (switches, hosts, or storage elements) and collects the appropriate information at fixed five-minute intervals.

Performance Manager uses a round-robin database to hold the statistical data collected from the fabric. This data is stored based on the configured parameters in the collection configuration file. At each polling interval, Performance Manager gathers the relevant statistics and stores them in the round-robin database. This database is a fixed size and will not grow beyond its preset limits.

To prevent loss of relevant information, Performance Manager creates a series of archived data to hold summarized information present in the real-time round-robin database. This archived data is used to generate daily, weekly, monthly, and yearly consolidated reports. In this way, Performance Manager maintains significant historical data without the cost of an ever-increasing database size.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Data Interpolation

One of the unique features of Performance Manager is its ability to interpolate data when statistical polling results are missing or delayed. Other performance tools may store the missing data point as zero, but this can distort historical trending. Performance Manager interpolates the missing data point by comparing the data point that preceded the missing data and the data point stored in the polling interval after the missing data. This maintains the continuity of the performance information.

## Data Collection

One year's worth of data for two variables (Rx and Tx bytes) requires a round-robin database (rrd) file size of 76K. If errors and discards are also collected, the rrd file size becomes 110K. The default internal values are:

- 600 samples of 5 minutes (2 days and 2 hours)
- 700 samples of 30 minutes (2 days and 2 hours, plus 12.5 days)
- 775 samples of 2 hours (above plus 50 days)
- 300 samples of 1 day (above plus 300 days, rounded up to 365)

A 1000-port SAN requires 110 MB for a year's worth of historical data that includes errors and discards. If there were 20 switches in this SAN with equal distribution of fabric ports, about two to three SNMP packets per switch would be sent every 5 minutes for a total of about 100 request or response SNMP packets required to monitor the data.

Flows, because of their variable counter requests, are more difficult to predict storage space requirements for. But as a rule of thumb, each extra flow adds another 76 kB.

The Performance Manager collector runs as a background process on the various supported operating systems. On Microsoft Windows, it runs as a service.

## Using Performance Thresholds

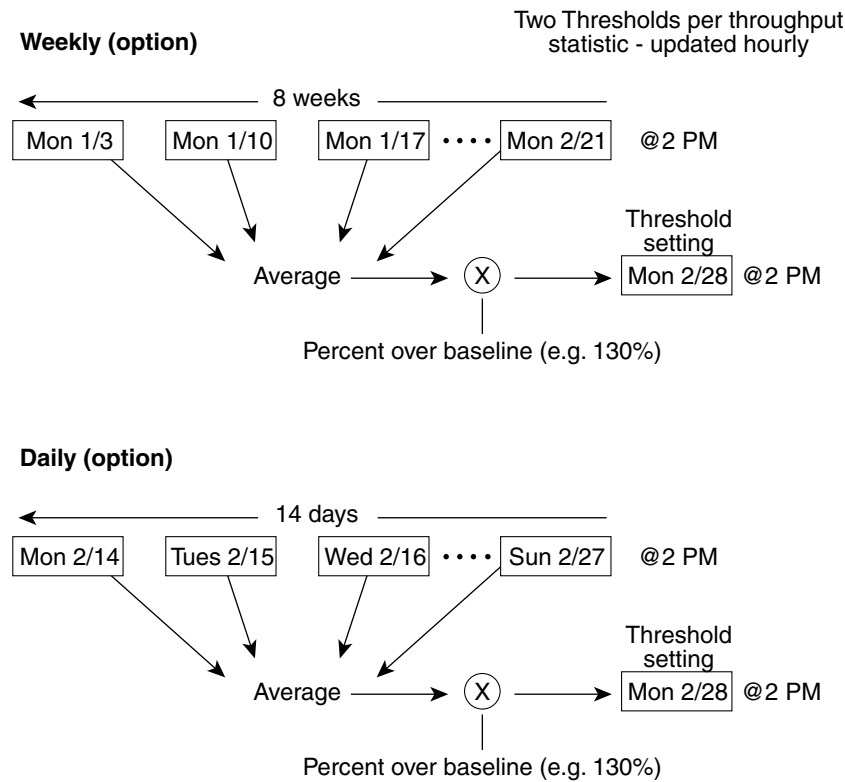
The Performance Manager Configuration Wizard allows you to set up two thresholds that will trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the Fabric Manager web client Events browser page.

Absolute value thresholds apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the Fabric Manager web client Events tab.

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every two weeks. Baseline thresholds are set as a percent of the weighted average (110% to 500%), where 100% equals the calculated weighted average. [Figure 6-1](#) shows an example of setting a baseline threshold for a weekly or daily option.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 6-1 Baseline Threshold Example**



The threshold is set for Monday at 2 PM. The baseline threshold is set at 130% of the weighted average for that statistic. The weighted average is calculated by averaging the statistics value that occurred at 2PM on Monday, for every prior Monday (for the weekly option) or the statistics value that occurred at 2PM on each day, for every prior day (for the daily option).

## Quick Data Collector and Flow Setup Wizards

The Performance Manager Flow and Performance Manager Setup wizards greatly simplify configuration. All you need to do is select the categories of statistics to capture and the wizards provide a list of flows and links to monitor. You can remove entries if desired, or just accept the provided list and start data collection. Statistics for host and storage links are not associated with a specific port on a switch, so you do not lose long term statistics if a connection is moved to a different port.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Authentication in Fabric Manager

---

Fabric Manager contains interdependent software components that communicate with the switches in your fabric. These components use varying methods to authenticate to other components and switches. This chapter describes these authentication steps and the best practices for setting up your fabric and components for authentication.

This chapter contains the following sections:

- [Fabric Manager Authentication Overview, page 7-1](#)
- [Best Practices for Discovering a Fabric, page 7-3](#)
- [Performance Manager Authentication, page 7-3](#)
- [Fabric Manager Web Services Authentication, page 7-4](#)

## Fabric Manager Authentication Overview


Fabric Manager contains multiple components that interact to manage a fabric. These components include:

- Fabric Manager client
- Fabric Manager server
- Performance Manager
- Interconnected fabric of Cisco MDS 9000 switches and storage devices
- AAA server (optional)

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

Figure 7-1 shows an example configuration for these components.

**Figure 7-1 Fabric Manager Authentication Example**



Administrators launch Fabric Manager client and select the seed switch that is used to discover the fabric. The username and password used are passed to Fabric Manager server and used to authenticate to the seed switch. If this username and password are not a recognized SNMP username and password, either Fabric Manager client or Fabric Manager server opens a CLI session to the switch (SSH or Telnet) and retries the username/password pair. If the username and password are recognized by the switch in either the local switch authentication database or through a remote AAA server, then the switch creates a temporary SNMP username that is used by Fabric Manager client and server.



**Note**

You may encounter a delay in authentication if you use a remote AAA server to authenticate Fabric Manager or Device Manager.



**Note**

You must allow CLI sessions to pass through any firewall that exists between Fabric Manager client and Fabric Manager server. See the [“Running Fabric Manager Behind a Firewall”](#) section on page 1-11.



**Note**

We recommend that you use the same password for the SNMPv3 username authentication and privacy passwords as well as the matching CLI username password.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Best Practices for Discovering a Fabric

Fabric Manager server monitors multiple physical fabrics under the same user interface. This facilitates managing redundant fabrics. A licensed Fabric Manager server maintains up-to-date discovery information on all configured fabrics so device status and interconnections are immediately available when you launch Fabric Manager client.

We recommend that you first discover the fabric through a username with network administrator or network operator roles. This ensures that Fabric Manager server has a complete view of the fabric. Subsequent Fabric Manager client sessions can filter this complete view based on the privileges of the client logging in. For example, if you have multiple VSANs in your fabric and you create users that are limited to a subset of these VSANs, you want to initiate a fabric discovery through Fabric Manager server using a network administrator or network operator role so that Fabric Manager server has a view of all the VSANs in the fabric. When a VSAN-limited user launches Fabric Manager client, that user sees only the VSANs they are allowed to manage.

## Setting up Discovery for a Fabric

To ensure that Fabric Manager server discovers your complete fabric, follow these steps:

- 
- Step 1** Create a special Fabric Manager administrative username in each switch on your fabric with network administrator or network operator roles. Or, create a special Fabric Manager administrative username in your AAA server and set every switch in your fabric to use this AAA server for authentication.
  - Step 2** Verify that the roles used by this Fabric Manager administrative username are the same on all switches in the fabric.
  - Step 3** Launch Fabric Manager client using the Fabric Manager administrative user. This ensures that your fabric discovery includes all VSANs.
  - Step 4** Set Fabric Manager Server to continuously monitor the fabric. See the [“Fabric Manager Server Fabric Monitoring and Removal” section on page 2-7](#).
  - Step 5** Repeat [Step 4](#) for each fabric that you want to manage through Fabric Manager server.
- 

## Performance Manager Authentication

Performance Manager uses the username and password information stored in the Fabric Manager server database. If this information changes on the switches in your fabric while Performance Manager is running, you need to update the Fabric Manager server database and restart Performance Manager. Updating the Fabric Manager server database requires removing the fabric from Fabric Manager server and rediscovering the fabric.

To update the username and password information used by Performance Manager, follow these steps:

- 
- Step 1** Choose **Server > Admin** in Fabric Manager. You see the Admin dialog box.
  - Step 2** Click the **Fabrics** tab to view the fabrics currently monitored by Fabric Manager server.
  - Step 3** Right-click the fabrics that have updated username and password information.
  - Step 4** Click **Remove** to remove these fabrics from Fabric Manager server.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Choose **File > Open Fabric**. You see the Open Fabric dialog box.
  - Step 6** Set the seed switch and the appropriate username and password to rediscover the fabric.
  - Step 7** Click **Open** to rediscover the fabric. Fabric Manager server updates its username and password information.
  - Step 8** Repeat [Step 5](#) through [Step 7](#) for any fabric that you need to rediscover.
  - Step 9** Click **Performance > Collector > Restart** to restart Performance Manager and use the new username and password.
- 

## Fabric Manager Web Services Authentication

Fabric Manager Web Services does not communicate directly with any switches in the fabric. Fabric Manager Web Services uses its own username and password combination that is either stored locally or stored remotely on an AAA server.

We recommend that you use a RADIUS server to authenticate users in Fabric Manager Web Services.

To configure Fabric Manager Web Services to use RADIUS authentication, follow these steps:

- 
- Step 1** Launch Fabric Manager Web Services. See the [“Launching and Using Fabric Manager Web Services” section on page 5-7](#).
  - Step 2** Choose **Admin > Web Users** to update the authentication used by Fabric Manager Web Services.
  - Step 3** Click **RADIUS**.
  - Step 4** Set the authentication.mode attribute to **radius**.
  - Step 5** Set the RADIUS server name, ports used, and shared secret for up to three RADIUS servers.
  - Step 6** Click **Modify** to save this information.
-



## Cisco Traffic Analyzer

---

Cisco Traffic Analyzer is a version of network top (ntop) software that is modified to support Fibre Channel and SCSI.

This chapter contains the following sections:

- [Using Cisco Traffic Analyzer with Performance Manager, page 8-1](#)
- [Installing Cisco Traffic Analyzer, page 8-4](#)
- [Configuring Performance Manager for Use with Cisco Traffic Analyzer, page 8-5](#)

## Using Cisco Traffic Analyzer with Performance Manager

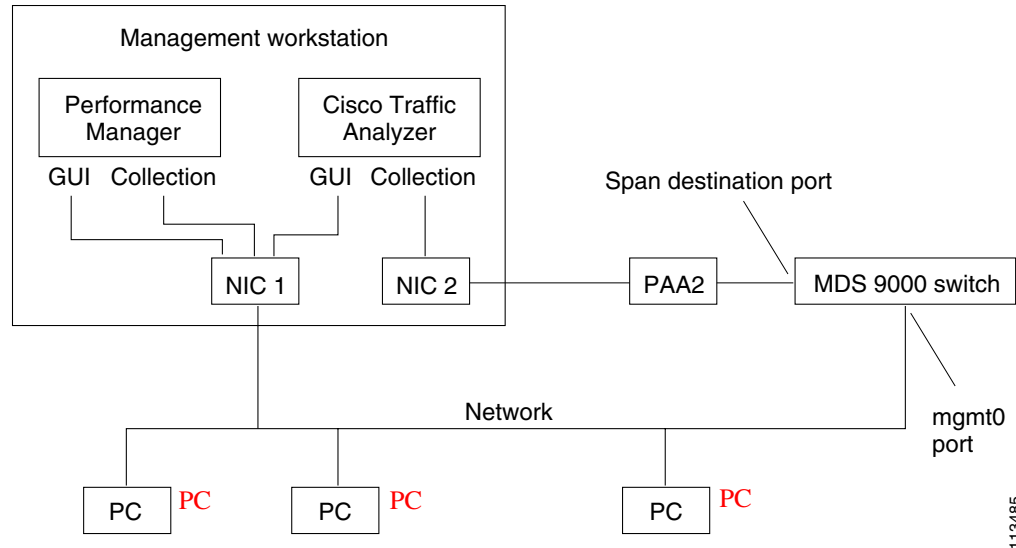
Performance Manager works in conjunction with Cisco Traffic Analyzer to monitor and manage the traffic on your fabric. Using Cisco Traffic Analyzer with Performance Manager requires the following components:

- A configured Fibre Channel Switched Port Analyzer (SPAN) destination (SD) port to forward Fibre Channel traffic.
- A Port Analyzer Adapter 2 (PAA-2) to convert the Fibre Channel traffic to Ethernet traffic.
- Cisco Traffic Analyzer software to analyze the traffic from the PAA-2.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

Figure 8-1 shows how Performance Manager works with Cisco Traffic Analyzer to monitor traffic on your fabric.

**Figure 8-1 Overview of Performance Manager Working with Cisco Traffic Analyzer**

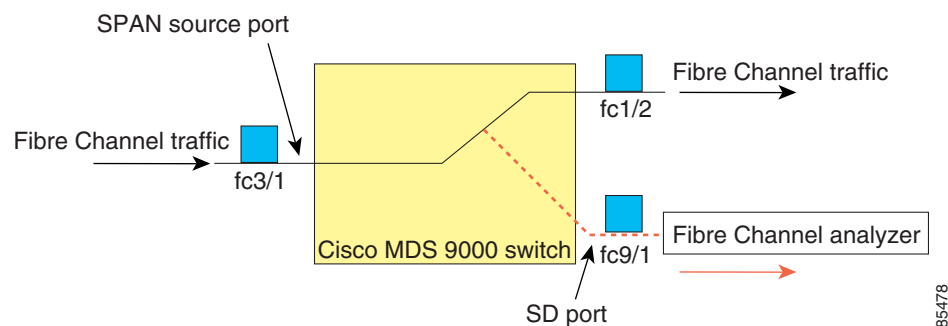


## Understanding SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel analyzer to the SD port to monitor SPAN traffic (see the “[Configuring World Wide Names](#)” section on page 23-3).

SD ports do not receive frames, they merely transmit a copy of the SPAN source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see [Figure 8-2](#)).

**Figure 8-2 SPAN Transmission**



For information on configuring SPAN, refer to the *Cisco MDS 9000 Family Configuration Guide*.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Understanding the PAA-2

The PAA-2 enables effective, low-cost analysis of Fibre Channel traffic. The device is a standalone Fibre Channel-to-Ethernet adapter, designed primarily to analyze SPAN traffic from a Fibre Channel port on a Cisco MDS 9000 Family switch. The main function of the Port Analyzer Adapter 2 is to encapsulate Fibre Channel frames into Ethernet frames. This allows low-cost analysis of Fibre Channel traffic while leveraging the existing Ethernet infrastructure.

The PAA-2 allows you to examine Fibre Channel frames of various sizes. Fibre Channel frames from Layers 2, 3, and 4 may be examined without network disruption.

## Understanding Cisco Traffic Analyzer

Performance Manager collects Fibre Channel level performance statistics using SNMP to access counters on Cisco MDS 9000 Family switches. To view detailed SCSI I/O statistics, you need to look at the data on an SD port with the help of Cisco Traffic Analyzer, which uses the Cisco Port Analyzer Adapter 2(PAA-2).

Cisco Traffic Analyzer provides real-time analysis of SPAN traffic or analysis of captured traffic through a Web browser user interface. Traffic encapsulated by one or more Port Analyzer Adapter 2 products can be analyzed concurrently with a single workstation running Cisco Traffic Analyzer, which is based on ntop, a public domain software enhanced by Cisco for Fibre Channel traffic analysis.

Round-trip response times, SCSI I/Os per second, SCSI read or traffic throughput and frame counts, SCSI session status, and management task information are monitored. Additional statistics are also available on Fibre Channel frame sizes and network management protocols.

For seamless performance analysis and troubleshooting, Cisco Traffic Analyzer can be launched in-context from Fabric Manager. Port World Wide Name (pWWN), Fibre Channel ID (FC ID), FC alias, and VSAN names are passed to Cisco Traffic Analyzer.

Cisco Traffic Analyzer must be downloaded and installed separately from the following website:

<http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml>.

Cisco Traffic Analyzer software is available under the Port Analyzer Adapter link. See the “[Installing Cisco Traffic Analyzer](#)” section on page 8-4.



### Caution

Cisco Traffic Analyzer for Fibre Channel throughput values are not accurate when used with the original Cisco Port Analyzer Adapter (PAA) if data truncation is enabled. PAA Version 2 (product ID DS-PAA\_2) is required to achieve accurate results with truncation, because it adds a count that enables Cisco Traffic Analyzer to determine how many data bytes were actually transferred.



### Note

Refer to the *Cisco MDS 9000 Family Configuration Guide* for information on configuring the settings for your span destination port. It is important that the data you collect through this port matches the data collected by Performance Manager through the mgmt0 port. If the data does not match, you cannot view Cisco Traffic Analyzer information through a Traffic Analyzer link on the detail page of a Performance Manager report.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

# Installing Cisco Traffic Analyzer

To install Cisco Traffic Analyzer on a UNIX workstation, follow these steps:

- 
- Step 1** Choose **Performance > Traffic Analyzer > Install** in Fabric Manager to access the web page where Cisco Traffic Analyzer is available.
  - Step 2** Download `fc-ntop.tar.gz` and install using the instructions it at the following website:  
<http://www.ntop.org>.
  - Step 3** Verify that the Fibre Channel port on the PAA-2 is connected to the SD port on the switch (Figure 8-1).
  - Step 4** Verify that the Ethernet port on the PAA-2 is connected to the workstation running Cisco Traffic Analyzer.
  - Step 5** Click **Interfaces > SPAN...** in Device Manager to configure SPAN on the required switch ports.
  - Step 6** Click **Interfaces > SPAN...** in Device Manager to verify that the Fibre Channel port connected to the PAA-2 is configured as an SD port. The port mode of the destination interface must be SD.
  - Step 7** Click the **Sessions** tab in Device Manager to verify the correct destination and source of traffic (ingress).
- 



## Caution

Cisco Traffic Analyzer must not be used with the PAA-2 in Management mode (MNM). Refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

To install Cisco Traffic Analyzer on a Windows workstation, follow these steps:

- 
- Step 1** Choose **Performance > Traffic Analyzer > Install** in Fabric Manager to access the web page where Cisco Traffic Analyzer is available.
  - Step 2** Download `ntop-win32.zip` and save it on your workstation.
  - Step 3** Unzip the downloaded file.



## Note

You need the WinPcap library file to use Cisco Traffic Analyzer on a Microsoft Windows system. You can download this file from the following website:  
<http://winpcap.polito.it/>.

- 
- Step 4** Open a command prompt and change directories to your ntop installation directory.
  - Step 5** Type `ntop -i` or install ntop as a service on Windows by following these steps:
    - a. Type `ntop /i` to install ntop as a service.
    - b. Choose **Start > Programs > Administrative Tools > Services** to access the Windows Services Panel.
    - c. Right-click `ntop` and choose **properties**. You see the Properties dialog box.
    - d. Set the Start Parameters to `-i <interface number>`, where *interface number* is the number of the interface on your workstation that connects to the PAA-2.
    - e. Click **Start** to start ntop on that interface.



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Note**

Subsequent restarts of the ntop service do not require setting the -i option, unless you are changing the interface that connects to the PAA-2.

- Step 6** Verify that the Fibre Channel port on the PAA-2 is connected to the SD port on the switch (Figure 8-1).
- Step 7** Verify that the Ethernet port on the PAA-2 is connected to the workstation running Cisco Traffic Analyzer.
- Step 8** Click **Interfaces > SPAN...** in Device Manager to configure SPAN on the required switch ports.
- Step 9** Click the **Sources** tab in Device Manager to verify that the Fibre Channel port connected to the PAA-2 is configured as an SD port. The port mode of the destination interface must be SD.
- Step 10** Click the **Sessions** tab in Device Manager to verify the correct destination and source of traffic (ingress).

**Tip**

To modify the script that launches ntop (ntop.sh or ntop.bat), follow the instructions provided within the script file. Create a backup of the original script before modifying the file.

—Linux platforms use the shell script path. The ntop output is sent to the syslog file (/var/log/messages by default).

—Windows platforms use the batch file. The ntop output is sent to a file located in the same directory as the one from which ntop is launched.

## Configuring Performance Manager for Use with Cisco Traffic Analyzer

To configure Performance Manager to work with Cisco Traffic Analyzer, follow these steps:

- Step 1** Get the following three pieces of information:
- The IP address of the management workstation on which you are running Performance Manager and Cisco Traffic Analyzer.
  - The path to the directory where Cisco Traffic Analyzer is installed.
  - The port that is used by Cisco Traffic Analyzer (the default is 3000).
- Step 2** Start Cisco Traffic Analyzer.
- a. Choose **Performance > Traffic Analyzer > Open**.
  - b. Enter the URL for Cisco Traffic Analyzer, in the format  
`http://<ip address>:<port number>`  
 where:  
`<ip address>` is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and `<port number>` is the port that is used by Cisco Traffic Analyzer (the default is :3000).
  - c. Click **OK**.
  - d. Choose **Performance > Traffic Analyzer > Start**.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

- e. Enter the location of Cisco Traffic Analyzer, in the format

D:\<directory>\ntop.bat

where:

D: is the drive letter for the disk drive where Cisco Traffic Analyzer is installed, and

<directory> is the directory containing the ntop.bat file.

- f. Click **OK**.

**Step 3** Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard.

**Step 4** Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard.

- a. Choose the VSAN you want to collect information for or choose All VSANs.

- b. Check the types of items you want to collect information for (hosts, ISLs, storage devices, and flows).

- c. Enter the URL for Cisco Traffic Analyzer in the format

http://<ip address>/<directory>

where:

<ip address> is the address of the management workstation on which you have installed Cisco Traffic Analyzer, and <directory> is the path to the directory where Cisco Traffic Analyzer is installed.

- d. Click **Next**.

- e. Review the end devices, and links you selected to make sure this is the data you want to collect.

- f. Click **Finish** to begin collecting data.



**Note**

Data is not collected for JBOD or for virtual ports. If you change the data collection configuration parameters during a data collection, you must stop and restart the collection process for your changes to take effect.

**Step 5** Click **Cisco Traffic Analyzer** at the top of the Host or Storage detail pages to view Cisco Traffic Analyzer information, or choose **Performance > Traffic Analyzer > Open**. Cisco Traffic Analyzer will not open unless ntop has been started already.



**Note**

For information on capturing a SPAN session and starting a Cisco Traffic Analyzer session to view it, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.



**Note**

For information on viewing and interpreting your Performance Manager data, see the [“Historical Performance Monitoring”](#) section on page 32-2.

For information on viewing and interpreting your Cisco Traffic Analyzer data, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

For performance drill-down, Fabric Manager Server can launch Cisco Traffic Analyzer in-context from the Performance Manager graphs. The aliases associated with hosts, storage devices, and VSANs are passed to Cisco Traffic Analyzer to provide consistent, easy identification.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 2**

# **Cisco MDS SAN-OS Installation and Configuration Files**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Obtaining and Installing Licenses

---

Licensing functionality is available in all switches in the Cisco MDS 9000 Family. This functionality allows you to access specified premium features on the switch after you install the appropriate license for that feature. Licenses are sold, supported, and enforced in Cisco MDS 9000 Family Cisco MDS SAN-OS Release 1.3(1) and later.

This chapter contains information related to licensing types, options, procedures, installation, and management for the Cisco MDS SAN-OS software.

This chapter includes the following sections:

- [Licensing Model, page 9-2](#)
- [Options to Install a License, page 9-5](#)
- [Installing Licenses Using Fabric Manager License Wizard, page 9-7](#)
- [Installing or Updating Licenses Using Device Manager, page 9-9](#)
- [Uninstalling Licenses, page 9-10](#)
- [Updating Licenses, page 9-11](#)
- [License Expiry Alerts, page 9-11](#)
- [Moving Licenses Between Switches, page 9-12](#)
- [Fabric Manager Server Licensing, page 9-12](#)

## Licensing Terminology

The following terms are used in this chapter:

- **Licensed feature**—Permission to use a particular feature through a license file, a hardware object, or a legal contract. This permission is limited to the number of users, number of instances, time span, and the implemented switch.
- **License expiry**—The time span during which a licensed feature is valid. The software tracks all licenses and sends periodic alerts before shutting down the licensed feature.
- **Counted license**—The number of usage instances for a licensed feature.
- **Licensed application**—A software feature that requires a license to be used.
- **License enforcement**—A mechanism that prevents a feature from being used without first obtaining a license.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Node-locked license—A license that can only be used on a particular switch using the switch's unique host ID.
- Host IDs—A unique chassis serial number that is specific to each Cisco MDS switch.
- Proof of purchase—A document entitling it's rightful owner to use licensed feature(s) on one Cisco MDS switch as described in that document. Also known as the claim certificate.
- Product Authorization Key (PAK)—The PAK allows you to obtain a license key from one of the sites listed in the proof of purchase document. After registering at the specified website, you will receive your license key file and installation instructions through e-mail.
- License key file—A switch-specific unique file that specifies the licensed features. Each file contains digital signatures to prevent tampering and modification. License keys are required to use a licensed feature. License keys are enforced within a specified time span.
  - License keys are required if your switch is running Cisco MDS SAN-OS Release 1.3 or later.
  - License keys are not enforced for features in Cisco MDS SAN-OS Release 1.2 or earlier.
- Counted license—The number of licenses issued for a single feature (for example, FCIP). You can increase counted licenses (incremental licenses) should a need arise in the future.
- Incremental license—An additional licensed feature that was not in the initial license file. License keys are incremental—if you purchase some features now and others later, the license file and the software detect the sum of all features for the specified switch.
- Evaluation license—A temporary license. Evaluation licenses are time bound (valid for a specified number of days) and are not tied to a host ID (switch serial number).
- Permanent license—A license that is not time bound (does not have an expiry date) is called a permanent license.
- Grace period—The amount of time the features in a license package can continue functioning without a license. The grace period is set to 120 days, and the countdown starts when either of the following two situations occur:
  - You are evaluating a feature for which you have not purchased a license.
  - You purchased a license that has reached its expiry date.

License packages can contain several features. If you disable a feature during the grace period, the clock does not stop for that feature, or for the other features in that license package. To suspend the grace period countdown for a licensed feature, you must disable every feature in that licence package.

- Support—If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

## Licensing Model

The licensing model defined for the Cisco MDS product line has two options:

- Feature-based licensing: features that are applicable to the entire switch. The cost varies based on a per-switch usage. [Table 9-1](#) lists the feature-based license packages.
- Module-based licensing: features that require additional hardware modules. An example is the IPS-8 or IPS-4 module using the FCIP feature.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 9-1 Feature-Based Licenses**

Feature License	Features
Standard package (free—no license required)	<ul style="list-style-type: none"> <li>• FCP, SSH, SFTP, and iSCSI protocols</li> <li>• Fabric Manager and Remote Monitoring (RMON)</li> <li>• VSANs, High availability, PortChannel, and Zoning</li> <li>• Fibre Channel Congestion Control (FCC)</li> <li>• Virtual Output Queuing (VOQ)</li> <li>• Diagnostics (SPAN, RSPAN, and FC Analyzer)</li> <li>• SNMP v3, Role-based access control, RADIUS</li> <li>• Call Home and Interoperability modes</li> <li>• IP access control lists (ACLs)</li> <li>• Terminal Access Controller Access Control System (TACACS+)</li> <li>• Fabric-Device Management Interface (FDMI)</li> <li>• Internet Storage Name Service (iSNS) client.</li> <li>• Dynamic Port VSAN Membership</li> <li>• Port Tracking</li> <li>• Distributed Device Alias Services</li> <li>• Cisco Fabric Services and PortChannel protocols</li> </ul>
Enterprise package (ENTERPRISE_PKG)	<ul style="list-style-type: none"> <li>• Enhanced security features: <ul style="list-style-type: none"> <li>– LUN zoning</li> <li>– Read-only zones</li> <li>– Port security</li> <li>– VSAN-based access control</li> <li>– Fibre Channel Security Protocol (FC-SP) authentication</li> <li>– IP Security Protocol (IPsec) using the MPS-14/2 module or the Cisco MDS 9216i Switch</li> </ul> </li> <li>• Advanced traffic engineering—Quality of Service (QoS)</li> <li>• Enhanced VSAN routing—inter-VSAN routing</li> <li>• Zone-based traffic prioritizing</li> <li>• Extended credits using the MPS-14/2 module or the Cisco MDS 9216i Multilayer Fabric Switch</li> <li>• Fibre Channel write acceleration</li> <li>• SCSI flow statistics</li> </ul>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 9-1 Feature-Based Licenses (continued)**

Feature License	Features
SAN extension over IP (SAN_EXTN_OVER_IP)	<p>SAN extension license are available for the IPS-8, IPS-4 and MPS-14/2.</p> <ul style="list-style-type: none"> <li>• FCIP protocol</li> <li>• FCIP compression</li> <li>• FCIP write acceleration</li> <li>• FCIP tape acceleration</li> <li>• SAN extension tuner features</li> <li>• IVR over FCIP</li> <li>• Hardware-based FCIP compression using the MPS-14/2 module or the Cisco MDS 9216i Multilayer Fabric Switch</li> </ul>
Mainframe (MAINFRAME_PKG)	<ul style="list-style-type: none"> <li>• FICON protocol and CUP management</li> <li>• FICON VSAN and intermixing</li> <li>• Switch cascading</li> <li>• Fabric binding</li> </ul>
Fabric Manager Server (FM_SERVER_PKG)	<ul style="list-style-type: none"> <li>• Multiple physical fabric management</li> <li>• Centralized fabric discovery services</li> <li>• Continuous MDS health and event monitoring</li> <li>• Long term historical Fibre Channel performance monitoring</li> <li>• Performance reports and charting for hotspot analysis</li> <li>• Fabric Manager Web Client for operational view</li> <li>• Threshold monitoring</li> <li>• Fabric Manager server proxy services</li> </ul>
Storage Services Enabler	<ul style="list-style-type: none"> <li>• Enables VERITAS Storage Foundation for Networks to run on the Advanced Services Module</li> </ul>

## Licensing High Availability

As with other Cisco MDS SAN-OS features, the licensing feature also maintains the following high availability standards for all switches in the Cisco MDS 9000 Family:

- Installing any license in any switch is a nondisruptive process.
- Installing a license automatically saves a copy of permanent licenses to the chassis in all switches.
- Enabling a license feature without a license key starts a counter on the grace period. You then have 120 days to install the appropriate license keys or disable the use of that feature. If at the end of the 120 day grace period the switch does not have a valid license key for the feature, the feature is automatically disabled by the switch.

Directors in the Cisco MDS 9500 Series have the following additional high availability features:

- The license software runs on both supervisor modules and provides failover protection.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- The license key file is mirrored on both supervisor modules. Even if both supervisor modules fail, the license file continues to function from the version that is available on the chassis.

## Options to Install a License

Licenses are installed on the switch using TFTP. If you have purchased a new switch through either your reseller or through Cisco Systems, you can:

- Obtain a factory-installed license (only applies to new switch orders).
- Perform a manual license Installation (applies to existing switches).

## Obtaining a Factory-Installed License

You can obtain factory-installed licenses for a new switch.

To obtain a factory-installed license for a new Cisco MDS switch, follow these steps.

---

**Step 1** Contact your reseller or Cisco representative and request this service.



**Note** If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

---

Your switch is shipped with the required licenses installed in the system. The proof of purchase document is sent along with the switch.

**Step 2** Obtain the host ID from the proof of purchase for future use.

**Step 3** Start to use the switch and the licensed features.

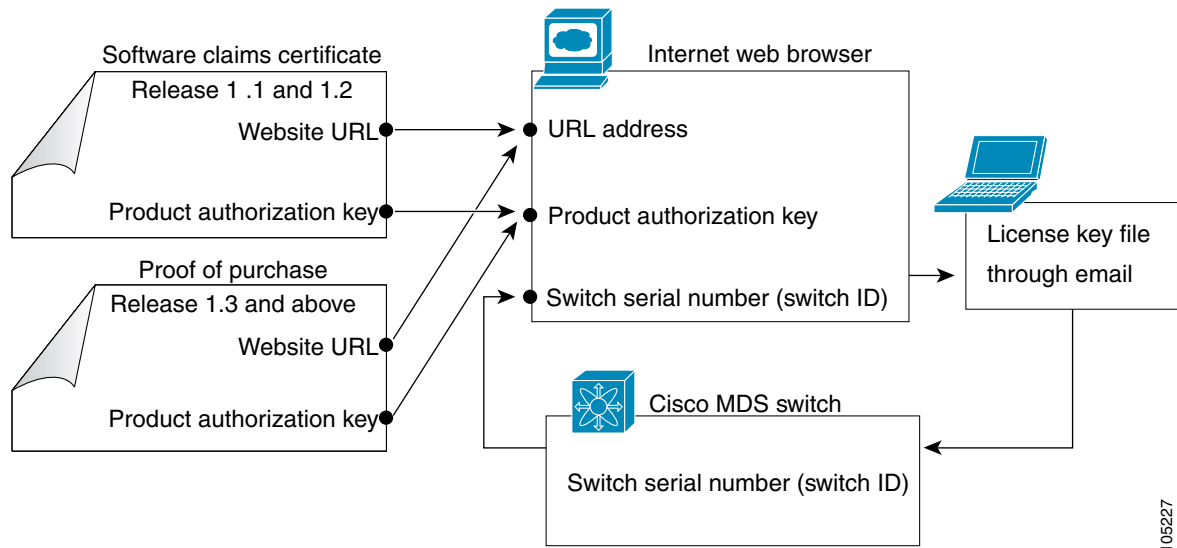
---

## Performing a Manual Installation

If you have existing switches or if you wish to install the licenses on your own, you must first obtain the license key file and then install that file in the switch (see [Figure 9-1](#)).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 9-1 Obtaining a License Key File**



105227

## Obtaining the License Key File

To obtain new or updated license key files, follow these steps.

- Step 1** Obtain the serial number for your switch. The host ID is also referred to as the switch serial number.
- Step 2** Obtain either your claim certificate or your proof of purchase document. This document accompanies every Cisco MDS switch.
- Step 3** Get the product authorization key (PAK) from either the claim certificate or the proof of purchase document.
- Step 4** Locate the website URL from either the claim certificate or the proof of purchase document.
- Step 5** Access the specified URL that applies to your switch and enter the switch serial number and the PAK. The license key file is sent to you by e-mail. The license key file is digitally signed to only authorize use on the requested switch. The requested features are also enabled once the Cisco MDS SAN-OS software on the specified switch accesses the license key file.



**Caution** Install the license key file in the specified MDS switch without making any modifications.

A license is either permanent or it expires on a fixed date. If you do not have a license, the grace period for using that license starts from the first time you start using a feature offered by that license (see the [“License Expiry Alerts” section on page 9-11](#)).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Installing the License Key File



### Tip

If you need to install multiple licenses in any switch in the Cisco MDS 9000 Family, be sure to provide unique file names for each license key file.

If you have purchased a new switch through either your reseller or through Cisco, you can have the licenses pre-installed in the factory, or you can install the licenses yourself. If you already have an existing switch, you install the licenses yourself. The best way to install licenses on the switches in your fabric is to use the License Wizard provided in Fabric Manager. You can also use Device Manager to install licenses on each switch individually.



### Note

You do not need a license to access a switch with Fabric Manager. See the [“Licensing Model” section on page 9-2](#) for a list of features requiring licenses.

You can install licenses two ways:


- Using the Fabric Manager License Wizard
- Using Device Manager

## Installing Licenses Using Fabric Manager License Wizard

To install licenses using the Fabric Manager License Wizard, follow these steps:

- Step 1** Log into a switch in the fabric containing the switches for which you want to install licenses.
- To install licenses on multiple switches, you do not need to log into each switch; however, the switches must be in the fabric you are viewing.
- Step 2** Start the License Wizard by clicking the **License Install Wizard** icon in the Fabric Manager toolbar (see [Figure 9-2](#)). Or you can select **Switches > Licenses** from the Physical Attributes pane. You see the license information in the Information pane, one line per feature. Click the **Keys** tab, and then click the **License Install Wizard** icon in the toolbar.

**Figure 9-2 License Install Wizard Icon**



You see the initial screen of the License Wizard.

- Step 3** If you have already obtained the license key files, click that radio button and proceed to Step 6.
- Step 4** Click **I have the Product Authorization Key (PAK)** if you have the authorization key.
- Step 5** Select the vendor from which you purchased your switch.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The License Server URL changes depending on the vendor you select. If your URL is different, or if you select **Other** as the vendor, enter the correct license server URL.



**Note** In some cases, license validation from Cisco partners requires Java version 1.4.2\_04 or higher. If you cannot install licenses from a Cisco partner, check to make sure your Java version is at least 1.4.2\_04.

- Step 6** Click **Next** to continue to the next screen.
- Step 7** Select the switches for which you have PAKs or license key files.
- When you check the checkbox for a switch, the PAK or license file name field for that switch becomes editable. The VDH=<serial number> for each switch is shown in the Host ID column.
- Step 8** Enter the PAK or license file name for each switch you have selected in the appropriate column. If you have the license files on your PC, you can double-click on the License File Name text area to bring up a dialog box and browse for the license files.
- You can install multiple licenses on the same switch using different PAKs. To do this, enter the PAKs separated by commas.
- Step 9** Click **Finish** to transfer the licenses from the host to the switches.
- Fabric Manager accesses the appropriate license site and installs the licenses onto each switch. The status of each installation is displayed in the Status column, as follows:
- success—Install or uninstall operation completed successfully.
  - inProgress—License install or uninstall operation is in progress.
  - corruptedLicenseFile—License file content is Invalid/Corrupted.
  - targetLicenseFileAlreadyExist—Target license file name already exists.
  - invalidLicenseFileName—License file does not exist.
  - duplicateLicense—License file is already installed.
  - generalLicensingFailure—General error from License Manager.
  - none—No install operation is performed.
  - licenseExpiryConflict—License exist with a different expiration date for the feature.
  - invalidLicenseCount—License count is invalid for the feature.
- Step 10** Click the **Close** button to close the wizard. To install more licenses at this point, you must close the wizard and launch it again.

## Viewing License Information in Fabric Manager

To view license information in Fabric Manager, follow these steps:

- Step 1** Select **Switches > Licenses** from the Physical Attributes pane. You see the license information in the Information pane, one line per feature.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 2** Click the **Feature Usage** tab to see the switch, name of the feature package, the type of license installed, the number of licenses used (Installed Count), the expiration date, the grace period (if you do not have a license for a particular feature), and any errors (for example, if you have a missing license). Click the **Keys** tab to display the information about each of the License Key files installed on your switches.



**Caution**

Once an expiration period has started, notifications about license expiration appear in the Fabric Manager's Events pane on a daily basis. During the last seven days of the expiration period, these messages are displayed hourly. After the final seven days of the expiration period, the feature is turned off and your network traffic may be disrupted.

- Step 3** Click the **Usage** tab to see the applications using the feature package on each switch. Use this tab to determine which applications depend on each license you have installed.

## Installing or Updating Licenses Using Device Manager

To install a license on your switch using Device Manager, follow these steps:

- Step 1** Select **Licenses** from the Admin menu.  
You see the Licenses dialog box.
- Step 2** Click the **Install** tab to display the Install fields.  
The HostId shows the "VDH=" portion of the serial number. The rest of the number is filled in when you complete Steps 3 through 5.
- Step 3** Enter the uniform resource identifier (URI) from which the license file will be retrieved for installation.  
You should already have copied the license file provided by Cisco.com or by some other means (for example, through the CLI) to this location.
- Step 4** Enter the target filename in the **Target Filename** field with which the license file will be installed.
- Step 5** Click **Install** if you are installing, or **Update** if you are updating.  
You see the status of the installation at the bottom of the dialog box, as follows:
- success—Install or uninstall operation completed successfully.
  - inProgress—License install/uninstall operation is in progress.
  - corruptedLicenseFile—License file content is Invalid/Corrupted.
  - targetLicenseFileAlreadyExist—Target license file name already exists.
  - invalidLicenseFileName—License file does not exist.
  - duplicateLicense—License file is already installed.
  - generalLicensingFailure—General error from license Manager.
  - none—No install operation is performed.
  - licenseExpiryConflict—License exist with a different expiration date for the feature.
  - invalidLicenseCount—License count is invalid for the feature.
  - notThisHost—License host-id in the license file doesn't match.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **licenseInGraceMore**—Number of licenses in grace period is more than the number in install license file.
- **licenseFileNotFound**—License file not found, for install/uninstall/update operation.
- **licenseFileMissing**—A previously installed license file is found missing.
- **invalidLicenseFileExtension**—License file does not have a .lic extension.
- **invalidURI**—Invalid license file URI, specified for install operation.
- **noDemoLicenseSupport**—Demo license not supported.
- **invalidPlatform**—Invalid Platform.

**Step 6** Repeat Steps 3 through 5 to install another license, or click **Close** to close the License Manager dialog box.

## Viewing License Information in Device Manager

To view license information in Device Manager, follow these steps:

- Step 1** Select **Licenses** from the Admin menu. You see the Licenses dialog box.
- Step 2** Click the **Features** tab to see the name of the feature package, the type of license, the expiration date, the grace period (if you do not have a license for a particular feature), and any errors, such as a missing license. Click the **Files** tab to display the information about each of the License Key files installed on your switch. Click the **Install** tab to install or update a license file. Click the **Usage** tab to see the applications that are using the features on the switch.

## Uninstalling Licenses

You can only uninstall a permanent license that is not in use. If you try to delete a permanent license that is currently being used, the software rejects the request and issues an error message. Uninstalling an unused license causes the grace period to come into effect. The grace period is counted from the first use of the feature without a license and is reset when a valid license file is installed.



### Note

Permanent licenses cannot be uninstalled if they are currently being used. Features turned on by permanent licenses must first be disabled, before that license is uninstalled.



### Tip

If you are using an evaluation license and would like to install a new permanent license, you can do so without service disruption and before the evaluation license expires. Removing an evaluation license immediately triggers a grace period without service disruption.



### Caution

Uninstalling a license requires the related features to first be disabled.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To uninstall the licenses, follow these steps:

- 
- Step 1** Log into the switch. If you are using Fabric Manager to remove licenses from multiple switches, you do not need to log in to each switch; however, the switches must be in the fabric you are viewing.
- Step 2** From Fabric Manager, select **Switches > Licenses** from the Physical Attributes pane. You see the license information in the Information pane, one line per feature.
- From Device Manager, select **Licenses** from the Admin menu. You see the Licenses dialog box.
- Step 3** In Fabric Manager, click the **Keys** tab. You see the list of License Key files. Click the name of the license you want to remove, and press the **Delete** key or click the **Delete Row** icon in the toolbar.
- In Device Manager, click **Uninstall**, and enter the name of the License Key file you want to remove. Click **Apply** to remove the License Key file, and click **Close** to close the dialog box.



**Note** To delete a license, you must disable the features enabled by that license. The delete procedure fails if the license is in use, and an error message is displayed.

---

## Updating Licenses

If your license is time bound, you must obtain and install an updated license. Contact technical support to request an updated license.



**Note** If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL:  
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

---

To update a license, follow these steps:

- 
- Step 1** Obtain the updated license file using the procedure described in the “[Obtaining the License Key File](#)” section on page 9-6.
- Step 2** Save your running configuration to a remote server.
- Step 3** Verify the name of the file to be updated.
- Step 4** Follow the procedure for updating a license, described in the “[Installing or Updating Licenses Using Device Manager](#)” section on page 9-9.
- 

## License Expiry Alerts

The Cisco MDS SAN-OS license counter keeps track of all licenses on a switch. Once an expiry period has started, you will receive console messages, SNMP traps, system messages, and Call Home messages on a daily basis.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Beyond that, the frequency of these messages become hourly during the last seven days of the expiry time span. The following example uses the FICON license feature.

Your FICON license feature is scheduled to expire in 60 days. If today is December 1st, the license expires on January 30th. In this case, you will receive:

- Daily alerts from December 1st to January 23rd.
- Hourly alerts from January 24th to January 29th.
- From January 30th, the FICON feature runs without a license for a grace period of 120 days.
- From January 30th to May 21st, you receive daily alerts about the grace period usage.
- From May 22nd to May 30th, you receive hourly alerts about the grace period ending.
- On May 31st, the FICON feature is automatically turned off.



#### Note

License expiry alerts cannot be configured.



#### Caution

After the final seven days of the grace period, the feature is turned off and your network traffic may be disrupted. The grace period also applies to licensed features in Cisco MDS 9000 Family Cisco MDS SAN-OS Release 1.2(x). While Cisco MDS SAN-OS Release 1.2(x) did not enforce licenses, any future upgrade will enforce license requirements and the 120-day grace period.

## Moving Licenses Between Switches

A license is specific to the switch for which it is issued and is not valid on any other switch. If you need to transfer a license from one switch to another, contact your customer service representative.



#### Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL:  
<http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml>

## Fabric Manager Server Licensing

When you install Fabric Manager, the a basic version of the Fabric Manager Server (FMServer) is installed with it. To get the enhanced features, such as Performance Manager and remote client support) you will need to buy and install the Cisco MDS 9000 Family Fabric Manager Server license package.

However, trial use of these enhanced features is available. To enable the 120-day trial, you simply use the feature as you would if you had purchased the license. You see a dialog box explaining that this is a demo version, enabled for a limited time.

If you are evaluating Fabric Manager Server features and want to stop the evaluation period for that feature, you can do that using Device Manager.

To stop the evaluation, follow these steps:

---

**Step 1** From Device Manager, select **Admin > Licenses**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You see the Licenses dialog box.

**Step 2** In the **Features** tab, select the feature you want to check in.

When you select the feature, you see a **Check In FM** button at the bottom of the dialog box.

**Step 3** Click **Check In FM** to stop the demo period timer.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Software Images

---

This chapter describes how to install and upgrade software images, and introduces the file system. It includes the following sections:

- [About Software Images, page 10-1](#)
- [Essential Upgrade Prerequisites, page 10-2](#)
- [Software Upgrade Methods, page 10-3](#)
- [Using the Software Install Wizard, page 10-4](#)
- [Upgrading from Cisco MDS SAN-OS 1.3\(4a\) to 2.0\(1b\), page 10-6](#)
- [File System Manipulation, page 10-7](#)

## About Software Images

Each switch is shipped with a Cisco MDS SAN-OS operating system for Cisco MDS 9000 Family switches. The Cisco MDS SAN-OS consists of two images—the kickstart image and the system image. To upgrade the switch to a new image, you must specify the variables which direct the switch to the images.

- To select the kickstart image use the KICKSTART variable.
- To select the system image use the SYSTEM variable.

The images and variables are important factors in any install procedure. You must specify the variable and the image to upgrade your switch. Both images are not always required for each install.

## Dependent Factors

The software image install procedure is dependent on the following factors:

- Software images—The kickstart and system image files reside in directories or folders that can be accessed from the Cisco MDS 9000 Family switch prompt.
- Image version—Each image file has a version.
- Flash disks on the switch—The bootflash: resides on the supervisor and the CompactFlash disk is inserted into the slot0: device.
- Supervisor modules—There are single or dual supervisor modules. In the dual supervisor scenario, the standby supervisor module should be updated first.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

# Essential Upgrade Prerequisites

Before attempting to migrate to any software image version, follow these guidelines:

- Customer Service

Before performing any software upgrade, contact your respective customer service representative to review your software upgrade requirements and to provide recommendations based on your current operating environment.



## Note

If you purchased Cisco support through a Cisco reseller, contact the reseller directly. If you purchased support directly from Cisco Systems, contact Cisco Technical Support at this URL: <http://www.Cisco.com/warp/public/687/Directory/DirTAC.shtml>

- Scheduling

Schedule the upgrade when the fabric is stable and steady. Ensure that everyone who has access to the switch or the network is not configuring the switch or the network during this time. All configurations are disallowed at this time.

- Space

Verify that sufficient space is available in the location where you are copying the images. This location includes the active and standby supervisor modules or bootflash: (internal to the switch).

- Standby supervisor module bootflash: directory (see the ).
- Internal bootflash offers approximately 200 MB of user space.

- Hardware

Avoid power interruption during any install procedure. These kinds of problems can corrupt the software image.

- Connectivity (to retrieve images from remote servers)

- Configure the IP address for the 10/100BASE-T Ethernet port connection (interface mgmt0).
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router to route traffic between subnets.
- Ensure any standby supervisor modules are physically connected with a path to the remote servers.

- Images

- Ensure that the specified system and kickstart images are compatible with each other.
- If the kickstart image is not specified, the switch uses the current running kickstart image.
- If you specify a different system image, ensure that it is compatible with the running kickstart image.
- Retrieve images in one of two ways:

Locally—images are locally available on the switch.

Remotely—images are in a remote location and the user specifies the destination using the remote server parameters and the file name to be used locally.

- Terminology

[Table 10-1](#) summarizes terms used in this chapter with specific reference to the install and upgrade process.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 10-1**      ***Terms Specific to This Chapter***

Term		Definition
bootable		The modules ability to boot or not boot based on image compatibility.
impact		The type of software upgrade mechanism—disruptive or nondisruptive.
install-type	reset	Resets the module.
	sw-reset	Resets the module immediately after switchover.
	rolling	Upgrades each module in sequence.
	copy-only	Updates the software for BIOS, loader, or bootrom.

## Software Upgrade Methods

You can upgrade software without any disruptions using the Cisco MDS SAN-OS software designed for mission-critical high availability environments. To realize the benefits of nondisruptive upgrades on the Cisco MDS 9500 Directors, we highly recommend that you install dual supervisor modules.

You can upgrade any switch in the Cisco MDS 9000 Family using one of the following methods:

- Automatic - you can use the Fabric Manager Software Install Wizard, as described in the [“Using the Software Install Wizard”](#) section on page 10-4.
- Manual - for information on manual upgrades, refer to the *Cisco MDS 9000 Family Configuration Guide*.

In some cases, regardless of which process you use, the software upgrades may be disruptive. These exception scenarios can occur under the following conditions:

- A single supervisor system with kickstart or system image changes.
- A dual supervisor system with incompatible system software images.

## Determining Compatibility

If the running image and the image you want to install are incompatible, the software reports the incompatibility. In some cases, you may decide to proceed with this installation. If the active and the standby supervisor modules run different versions of the image, both images may be HA compatible in some cases and incompatible in others.

Compatibility is established based on the image and configuration:

- Image incompatibility—The running image and the image to be installed are not compatible.
- Configuration incompatibility—There is a possible compatibility if certain features in the running image are turned off as they are not supported in the image to be installed. The image to be installed is considered incompatible with the running image if one of the following statements is true:
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and may cause the switch to move into an inconsistent state. In this case, the incompatibility is *strict*.
  - An incompatible feature is enabled in the image to be installed and it is not available in the running image and does not cause the switch to move into an inconsistent state. In this case, the incompatibility is *loose*.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Recognizing Failure Cases

The following situations will cause the installation process to end:

- If the standby supervisor module bootflash: directory does not have sufficient space to accept the updated image.
- If the specified system and kickstart images are not compatible.
- If the installation procedure is performed on the standby supervisor module.
- If the fabric or switch is configured while the upgrade is in progress.
- If a module is removed while the upgrade is in progress.
- If the switch has any power disruption while the upgrade is in progress.
- If the entire path for the remote location is not specified accurately.
- If images are incompatible after an upgrade. For example, a switching module image may be incompatible with the system image, or a kickstart image may be incompatible with a system image.



### Caution

Avoid ending the switch progress after starting the installation process. If the installation process has ended, be sure to verify the state of the switch at every stage, and wait 10 second before attempting to restart the installation process. If you restart the installation process before waiting 10 seconds, the process will not start and you will see an error message indicating that an installation is currently in progress.

## Using the Software Install Wizard



### Note

Before you use this wizard, be sure the standby supervisor management port is connected.

To use the Software Install Wizard, follow these steps:

- Step 1** Open the Software Install Wizard by clicking on its icon in the toolbar (see [Figure 10-1](#)).

**Figure 10-1** Software Install Wizard Icon



You see the Software Install Wizard.

- Step 2** Select the switches you want to install images on. You must select at least one switch in order to proceed. When finished, click **Next**.
- Step 3** Optionally, check the **Skip Image Download** check box and click **Next** to use images that are already downloaded (the file is already on the bootflash). Proceed to [Step 7](#).



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Click on the row under the System, Kickstart, Asm-sfn, or ssi columns to enter image URIs. You must specify at least one image for each switch in order to proceed.
- Step 5** Check the active (and standby, if applicable) bootflash on each switch to see if there is enough space for the new images. You can see this information in the Flash Space column.
- This screen shows the active (and standby, if applicable) bootflash space on each switch, and shows the status (whether there is enough space for the new images). If any switch has insufficient space, you cannot proceed. Deselect the switch without enough bootflash by going back to the first screen and unchecking the check box for that switch.
- Step 6** Click **Next**. The Select Download Image page displays.
- Step 7** Double-click the table cell under System, Kickstart, Asm-sfn, or Ssi and select from a dropdown list of images available in bootflash on each switch. You must select at least one image for each switch to proceed.



**Note** There is no limit on the number of switches you can upgrade. However, the upgrade is a serial process; that is, only a single switch is upgraded at a time.

- Step 8** Click **Next**. The final verification page displays.
- Step 9** Optionally, check **Ignore version check results** to bypass a version check.



**Note** The version check provides information about the impact of the upgrade for each module on the switch. It also shows any HA-related incompatibilities that might result. You see a final dialog box at this stage, prompting you to confirm that this check should be performed. We recommend that you do not ignore the version check results.



**Caution** If **Ignore version check results** is checked, the upgrade will proceed even if the current switch version is newer than the version you are installing.

- Step 10** Click **Finish** to start the installation or click **Cancel** to leave the installation wizard without installing new images.



**Note** On hosts where the TFTP server cannot be started, a warning is displayed. The TFTP server may not start because an existing TFTP server is running or because access to the TFTP port 69 has been denied for security reasons (the default setting on linux). In these cases, you cannot transfer files from the local host to the switch.



**Note** Before exiting the session, be sure the upgrade process is complete. The wizard will display a status as it goes along. Check the lower left-hand corner of the wizard for the status message **Upgrade Finished**. First, the wizard displays the message **Success** followed a few seconds later by **InProgress Polling**. Then the wizard displays a second message **Success** before displaying the final **Upgrade Finished**.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Upgrading from Cisco MDS SAN-OS 1.3(4a) to 2.0(1b)

To upgrade a switch from 1.3(4a) to 2.0(1b), use Device Manager to copy the image files to bootflash and then use FM to perform the upgrade procedure. This procedure assumes you are using Device Manager 1.3(4a) or higher.

To copy the image files from a server or PC to bootflash, follow these steps:

- 
- Step 1** Start TFTP, FTP, SCP, or SFTP on the server or PC where you have the image files stored.
  - Step 2** In Device Manager, select **Admin > Flash Files**. You see the bootflash directory listed for the supervisor's local partition, by default.
  - Step 3** Select the device and partition from the dropdown lists for the directory containing the file you want to copy.
  - Step 4** Click the **Copy** button to open the Copy dialog box.
  - Step 5** Select the protocol you want to use to perform the copy procedure.
  - Step 6** Enter the address of the source server.
  - Step 7** If necessary, enter your remote username and password on that server.
  - Step 8** Click the ... button after the SourceName field to browse for the source file on your local PC or on the server, depending on the type of copy.
  - Step 9** Enter the destination name for the file.



**Note** If you are copying to Flash, the file name must be of the form [device>:][<partition>:]<file>

where <device> is a value obtained from FlashDeviceName,  
<partition> is obtained from FlashPartitionName  
and <file> is any character string that does not have embedded colon characters.

---

- Step 10** Click **Apply**.
- 

To upgrade with Fabric Manager, you use the Software Install Wizard. Software upgrades may be disruptive under the following conditions:

- A single supervisor system with kickstart or system image changes.
- A dual supervisor system with incompatible system software images.



**Note** Before you use the Software Install Wizard, be sure the standby supervisor management port is connected.


---

To use the Software Install Wizard, follow these steps:

- 
- Step 1** Open the Software Install Wizard by clicking on its icon in the toolbar (see [Figure 10-1](#)).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 10-2 Software Install Wizard Icon**



You see the Software Install Wizard.

**Step 2** Select the switches from the list shown, for which you want to upgrade or install images.

You must select at least one switch in order to proceed. When finished, click **Next**.

**Step 3** Specify the new images to use for each switch model.

To use images that are already downloaded (the file is already on the bootflash), check the **Skip Image Download** checkbox.

**Step 4** Double-click the table cell under System, Kickstart, or Asm-sfn to see a dropdown list of images to choose from.

**Step 5** Select an image to use for the upgrade.

You must select at least one image for each switch to proceed.



**Note**

There is no limit on the number of switches you can upgrade. However, the upgrade is a serial process; that is, only a single switch is upgraded at a time.

**Step 6** Start the upgrade.

If you check **version check** before the upgrade process is started, a version check is done. This check provides information about the impact of the upgrade for each module on the switch. It also shows any HA-related incompatibilities that might result. You see a final dialog box at this stage, prompting you to confirm that this check should be performed.



**Caution**

If **version check** is enabled, the upgrade will proceed even if your version is newer than the version you are installing.



**Note**

Before exiting the session, be sure the upgrade process is complete. The wizard will display a status as it goes along. Check the lower left-hand corner of the wizard for the status message **Upgrade Finished**. First, the wizard displays the message **Success** followed a few seconds later by **InProgress Polling**. Then the wizard displays a second message **Success** before displaying the final **Upgrade Finished**.

## File System Manipulation

All switches in the Cisco MDS 9000 Family have one internal bootflash: that resides in the supervisor or switching module. You have access to two directories within the internal bootflash: file system.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- The volatile: directory provides temporary storage, and it is also the default. Files in temporary storage (volatile:) are erased when the switch reboots.
- The bootflash: (nonvolatile storage) directory provides permanent storage. Files in permanent storage (bootflash:) are preserved through reboots and power outages.

Cisco MDS 9500 Series directors contain an additional external CompactFlash referred to as the slot0: directory. The external CompactFlash, an optional device for MDS 9500 Series directors, can be used for storing software images, logs, and core dumps.

You can use Device Manager to perform the following functions to help you manage software image files and configuration files:

- [Listing the Files in a Directory, page 10-8](#)
- [Creating a Directory, page 10-8](#)
- [Deleting an Existing File or Directory, page 10-9](#)
- [Copying Files, page 10-9](#)
- [Performing Other File Manipulation Tasks, page 10-10](#)

## Listing the Files in a Directory

To list the files in a directory using Device Manager, follow these steps:

- 
- Step 1** Select **Admin > Flash Files**. By default, you see the bootflash: directory listed for the supervisor's local partition.
- Step 2** Select the device and partition from the drop-down lists for the directory you want to view. You see a list of files and directories.
- 

## Creating a Directory

To create a directory using Device Manager, follow these steps:

- 
- Step 1** Select **Admin > Flash Files**. By default, you see the bootflash: directory listed for the supervisor's local partition.
- Step 2** Select the device and partition from the drop-down lists for the directory where you want to create the new directory.
- Step 3** Click the ... button to create a new directory. You see the Create New Directory dialog box.
- Step 4** Enter the name of the new directory, and click **OK**. You see the new directory in the directory listing.




---

**Tip** Any directory saved in the volatile: file system is erased when the switch reboots.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Deleting an Existing File or Directory

To delete a file or directory using Device Manager, follow these steps:

- 
- Step 1** Select **Admin > Flash Files**. By default, you see the bootflash: directory listed for the supervisor's local partition.
  - Step 2** Select the device and partition from the drop-down lists for the directory containing the file or directory you want to delete.
  - Step 3** Click to select the file or directory you want to delete.
  - Step 4** Click the **Delete** button to delete the file or directory.

**Caution**

If you specify a directory, the **delete** command deletes the entire directory and all of its contents.

---

## Copying Files

To copy a file using Device Manager, follow these steps:

- 
- Step 1** Select **Admin > Flash Files**. By default, you see the bootflash: directory listed for the supervisor's local partition.
  - Step 2** Select the device and partition from the drop-down lists for the directory containing the file you want to copy.
  - Step 3** Click the **Copy** button. You see the Copy dialog box.
  - Step 4** Select the protocol you want to use to perform the copy procedure.
  - Step 5** Enter the address of the source server (Flash to Flash copy only).
  - Step 6** Click the ... button to browse for the source file on your local PC or on the server, depending on the type of copy.

**Note**

If you are copying from Flash, the file name must be of the form  
[device>:][<partition>:]<file>

where <device> is a value obtained from FlashDeviceName,  
<partition> is a value obtained from FlashPartitionName and  
<file> is the name of a file in Flash.

---

- Step 7** Enter the destination name for the file.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

If you are copying to Flash, the file name must be of the form  
[device>:][<partition>:]<file>

where <device> is a value obtained from FlashDeviceName,  
<partition> is a value obtained from FlashPartitionName and  
<file> is any character string that does not have embedded colon characters.

**Step 8** Click **Apply**.

---

## Performing Other File Manipulation Tasks

To perform the following CLI-specific tasks, refer to the *Cisco MDS 9000 Family Configuration Guide*:

- Displaying file contents
- Displaying the last line in a file
- Saving output to a file
- Moving files
- Compressing and uncompressing files
- Executing commands specified in a script
- Setting the delay time



## CHAPTER 11

# Configuration Files

---

This chapter describes how to update configuration files. It includes the following sections:

- [Working with Configuration Files, page 11-1](#)
- [Saving the Configuration File, page 11-1](#)
- [Copying the Configuration File, page 11-2](#)

## Working with Configuration Files

Configuration files can contain some or all of the commands needed to configure one or more switches. For example, you might want to download the same configuration file to several switches that have the same hardware configuration so that they have identical module and port configurations.

You can configure a switch in the Cisco MDS 9000 Family by using configuration files you create or download from another switch. In addition, you can store configuration files on a bootflash device on the supervisor module and you can configure the switch using a configuration stored on an external CompactFlash disk. Before you begin downloading a configuration file using a remote server, do the following:

- Ensure the configuration file to be downloaded is in the correct directory on the remote server.
- Ensure that the permissions on the file are set correctly. Permissions on the file should be set to world-read.
- Ensure the switch has a route to the remote server. The switch and the remote server must be in the same subnetwork if you do not have a router or default gateway to route traffic between subnets.
- Check connectivity to the remote server using the **ping** command.

## Saving the Configuration File

To copy the configuration file, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | In Device Manager, select <b>Admin &gt; Save Configuration</b> .<br>You see the message “Really save running to startup configuration?” |
| <b>Step 2</b> | Click <b>Yes</b> to save the configuration. Click <b>No</b> to close the popup without saving the configuration.                        |
-

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Copying the Configuration File

You can copy the running configuration to the startup configuration across the entire fabric in Cisco MDS SAN-OS Release 2.1(1a) or later by using the `fabricStartupConfig` option. This triggers every switch in the fabric to copy its running configuration to its startup configuration.



### Note

If any switch fails during this fabric-wide copy, that switch and the switch that you used to initiate this command will keep the existing startup configuration. This does not affect the other switches in the fabric.

To copy the configuration file, follow these steps:

- 
- Step 1** In Device Manager, select **Admin > Copy Configuration**. You see the Copy Configuration dialog box.
  - Step 2** Select the location of the file you want to copy From (serverFile, startupConfig, runningConfig).
  - Step 3** Select the location of the file you want to copy To (serverFile, runningConfig, fabricStartupConfig)



### Note

You can copy a file fabric-wide using the `fabricStartupConfig` option, available in Cisco MDS SAN-OS Release 2.1(1a) or later.

- 
- Step 4** Enter the address of the source server.
  - Step 5** Click the ... button to browse for the source file on the switch or the server, depending on the type of copy.
  - Step 6** Select the protocol you want to use to perform the copy procedure.
  - Step 7** Enter the username and password you use to access the switch or server.
  - Step 8** Click **Apply**.
-



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 3**

### **Switch Configuration**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Cisco Fabric Services

---

This chapter contains descriptions of, and instructions for using, the Cisco Fabric Services (CFS) with Fabric Manager and Device Manager. The Cisco MDS SAN-OS software uses CFS to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

Several Cisco MDS SAN-OS features use the CFS infrastructure to maintain and distribute the contents of a particular feature's database.

This chapter contains the following sections:

- [About CFS, page 12-1](#)
- [Enabling CFS for a Feature, page 12-3](#)
- [Disabling or Enabling CFS Distribution on a Switch, page 12-6](#)
- [Disabling or Enabling CFS Distribution on a Switch, page 12-6](#)
- [A CFS Example Using Fabric Manager, page 12-7](#)
- [A CFS Example Using Device Manager, page 12-9](#)

## About CFS

Many features in the Cisco MDS 9000 Family of switches require configuration synchronization in all switches in the fabric. Maintaining configuration synchronization across a fabric is important to maintain fabric consistency. As of Cisco MDS SAN-OS Release 2.0, Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the features. CFS can discover CFS capable switches in the fabric and discover feature capabilities in all CFS capable switches.

## Cisco MDS SAN-OS Features Using CFS

The following Cisco MDS SAN-OS features use the CFS infrastructure:

- TACACS+
- RADIUS
- Role-based access
- Dynamic port VSAN membership

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Distributed Device Alias Services
- iSNS
- Call Home
- VSANs
- IVR topology
- Port security
- NTP
- Syslog
- SCSI flow services
- fctimer
- Fabric Start Config Manager

**Note**

Most tabs in the Information pane for features using CFS are dimmed until you visit the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is visited, the other tabs activated.

## CFS Features

CFS has the following features:

- Two modes of distribution:
  - Uncoordinated distributions where multiple parallel distributions are allowed in the fabric.
  - Coordinated distributions where only one distribution is allowed in the fabric at any given time.
- Two scopes of distribution:
  - Logical scope where the distribution happens within the scope of a VSAN. An example feature is port security, where the configuration database is applicable only within a VSAN.
  - Physical scope where the distribution span the entire physical topology. Example features are NTP and DPVM (WWN based VSAN), which are independent of VSANs.
- A peer-to-peer protocol that does not have a client-server relationship at the CFS layer.
- A merge protocol that facilitates the merge of feature configuration during a fabric merge event (when two independent fabrics merge).

All switches in the fabric must be CFS capable. Non-CFS capable switches do not receive the distribution and may prevent part of the fabric from receiving the intended distribution.

CFS has the following attributes:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled feature, the configuration modification process begins and the feature locks the fabric.
- Temporary buffer—This term differs based on the feature:
  - Some feature saves the changes to the temporary buffer when you perform configuration modifications. The changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the contents of the effective database are overwritten by the contents of the temporary buffer.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Other features make a copy of the configuration database (when you start making configuration changes) and use this as a working copy. When you commit the changes, the working copy becomes the new configuration database. In this case, this working copy is the temporary buffer.
- Explicit commit—Each feature requires an explicit **commit** action to commit the changes in the temporary buffer. The changes in the temporary buffer are not applied if you do not perform the commit. The **commit** action distributes the new database in the fabric and then releases the fabric lock.
- CFS distribution is enabled or disabled on a per-feature basis. The default (enable or disable) for each feature differs between features. If CFS distribution is disabled for a feature, then that feature does not distribute any configuration nor does it accept a distribution from other switches in the fabric.

## Enabling CFS for a Feature

All CFS based features provide an option to enable or disable the distribution capabilities. Features that existed prior to Cisco MDS SAN-OS Release 2.x have the distribution capability disabled by default and must have their distribution capabilities enabled explicitly.

Features that are introduced in Cisco MDS SAN-OS Release 2.x have the distribution enabled by default but can be disabled as required.

The feature configuration is not distributed by CFS unless distribution is explicitly enabled for that feature.

To enable CFS for a feature using Fabric Manager, follow these steps:

---

**Step 1** Choose the feature that you want to enable CFS for. For example, choose **Switches > Events > CallHome** from the Physical Attributes pane. The Information pane shows that feature, with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

**Step 2** Decide which switch(es) to enable CFS for. Set the Admin state column to either **enable** to enable CFS, or **disable** to disable CFS.




---

**Note** You should enable CFS for all switches in the fabric or VSAN for the feature that uses CFS.


---

**Step 3** Right-click anywhere on the row to get the pop-up menu. Select **Apply Changes** to apply the CFS configuration change. The CFS tab updates as the CFS changes take affect. Fabric Manager retrieves the status of the CFS change and updates the Last > Result column.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To enable CFS for a feature using Device Manager, follow these steps:

- 
- Step 1** Select **Admin > CFS** from the main menu. You see the CFS dialog box with the CFS status for all features on that switch.
- Step 2** Decide which feature(s) to enable CFS for. Set the Command column to either **enable** to enable CFS, or **disable** to disable CFS.
- 
-  **Note** You should enable CFS for all switches in the fabric or VSAN for the feature that uses CFS.
- 
- Step 3** Select the **Pending Differences** button to check the configuration of this feature on this switch as compared to other switches in the fabric or VSAN that have CFS enabled for this feature.
- Step 4** Select **Apply** to apply the CFS configuration change or **Close** to close the dialog box without making changes. Device Manager retrieves the status of the CFS change and updates the Last > Result column.
- 

## Locking the Fabric

When you configure (first time configuration) a CFS-enabled feature, that feature starts a CFS session and locks the fabric. When a fabric is locked, the Cisco MDS SAN-OS software does not allow any configuration changes from a switch (other than the switch holding the lock) to this Cisco MDS SAN-OS feature, and issues a message to inform you about the locked status. The configuration changes are held in a pending database by that feature.

If you start a CFS session that requires a fabric lock but forget to end the session, an administrator can clear the session. If you lock a fabric at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

## Committing Changes

A commit causes the pending database to be saved by all feature peers. It also causes locks to be released at all switches.

In general, the commit function does not start a session; only a lock function starts a session. However, an empty commit is allowed if configuration changes were not previously made. In this case, a commit results in a session whereby locks are acquired and the current database distributed.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The feature applies the changes locally and releases the fabric lock.
- None of the external switches report a successful state—The feature considers this state a failure and does not apply the changes to any switch in the fabric. The fabric lock is not released.

Fabric locks are released even if all peers failed to commit the configuration.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To commit changes using Fabric Manager for CFS-enabled features, follow these steps:

- 
- Step 1** Choose the feature you want to enable CFS for. For example, choose **Switches > Events > CallHome** from the Physical Attributes pane. The Information pane shows that feature, with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.
- Step 2** Select **Config Changes** for any switch. Set the Action column to **commit**.
- Step 3** Click the **Apply Changes** to commit the configuration changes for that feature and distribute the changes through CFS, or click **Undo Changes** to discard the changes for that feature. Fabric Manager retrieves the status of the CFS change and updates the Last > Result column on a per feature or per VSAN basis, depending on the feature.
- 

To commit changes using Device Manager for CFS-enabled features, follow these steps:

- 
- Step 1** Select **Admin > CFS** from the main menu. You see the CFS dialog box with the CFS status for all features on that switch.
- Step 2** For each applicable feature, set the Command column to **commit** to commit the configuration changes for that feature and distribute the changes through CFS, or set it to **abort** to discard the changes for that feature and release the fabric lock for CFS for that feature.
- Step 3** Optionally, select **none** or **VsanID** as the basis for the CFS distribution for CFS features that require this.
- Step 4** Select the **Pending Differences** button to check the configuration of this feature on this switch as compared to other switches in the fabric or VSAN that have CFS enabled for this feature.
- Step 5** Select **Apply** to apply the CFS configuration change or **Close** to close the dialog box without making changes. Device Manager retrieves the status of the CFS change and updates the Last > Result column.
- 



**Caution**

If you do not commit the changes, they are not saved to the running configuration.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Clearing a Locked Session

You can clear locks held by a feature from any switch in the fabric. This option is provided to rescue you from situations where locks are acquired and not released. This function requires administrative permissions.



### Caution

Exercise caution when using this function to clear locks in the fabric. Any pending configurations in any switch in the fabric is flushed and lost.

## Disabling or Enabling CFS Distribution on a Switch

By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation.

As of Cisco MDS SAN-OS Release 2.1(1a), you can globally disable CFS on a switch to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

To globally disable or enable CFS distribution on a switch using Fabric Manager, follow these steps:

- Step 1** Choose any CFS feature. For example, choose **Switches > Events > CallHome** from the Physical Attributes pane. The Information pane shows that feature, with a CFS tab.
- Step 2** Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.
- Step 3** Set the **Enable > Global** drop-down menu to **disable** or **enable** for all switches that you want to disable or enable CFS on.
- Step 4** Set the **Config Changes > Action** column to **commit**.
- Step 5** Click the **Apply Changes** to commit the configuration changes for that feature and distribute the changes through CFS, or click **Undo Changes** to discard the changes for that feature.

To globally disable or enable CFS distribution on a switch using Device Manager, follow these steps:

- Step 1** Select **Admin > CFS** from the main menu. You see the CFS dialog box with the CFS status for all features on that switch.
- Step 2** Uncheck or check the **Globally Enabled** check box to disable or enable CFS distribution on this switch.
- Step 3** Select **Apply** to disable CFS on this switch or **Close** to close the dialog box without making changes.



*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## CFS Merge Support

A feature keeps the configuration synchronized in a fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification each time a feature peer comes online. If two fabrics with M and N feature peers merge and if a feature triggers a merge action on every such notification, a link up event results in M\*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per feature per scope. The protocol involves selecting one switch in a fabric as the merge master for that fabric. The other switches do not play any role in the merge process.

During a merge, the merge master in the two fabrics exchange their configuration databases with each other. The feature on one of the databases merges the information, decides if the merge is successful or not, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database. The merge master in a fabric is the master or seed switch that is used for distributing each CFS-enabled feature.

## A CFS Example Using Fabric Manager

This procedure is an example of what you see when you use Fabric Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature's documentation.

To configure a feature that uses CFS using Fabric Manager, follow these steps:

- Step 1** Select the CFS-capable feature you want to configure. For example, Choose **Switches > Clock > NTP** from the Physical Attributes pane. You see the feature configuration in the Information pane, including a CFS tab.
- Step 2** Click the **CFS** tab. You see the CFS configuration and status for each switch.
- Step 3** Choose **enable** from the Enable > Admin drop-down box for all switches in the fabric



**Note** A warning displays if you do not enable CFS for all switches in the fabric for this feature.

- Step 4** Check the **Master** check box for the switch that you want to act as the merge master for this feature.
- Step 5** Choose **commit** from the Config Changes > Action drop-down box for each switch that you enabled CFS on.
- Step 6** Choose a configuration tab. For example, choose **Servers** from the Information pane for NTP. You see the configuration for this feature based on the master switch.
- Step 7** Modify the feature configuration. For example, click on **Create Row** to create a new server for NTP.
  - a. Set the **ID**, and the **Name or IP Address** for the NTP server.
  - b. Set the **Mode** radio button and click **Add** to add the server.  
Fabric Manager sends the request to the master switch and updates the CFS > Last > Results column with the "pending" status.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 8** Choose the **CFS** tab for this feature.

**Step 9** Set the Config Changes > Action column to **commit** to distribute the feature change through the fabric. Fabric Manager only changes the status to “running” when **commit**, **clear**, or **abort** is selected and applied.



**Note** Fabric Manager will not change the status to “pending” if **enable** is selected, because the “pending” status does not apply until the first actual change is made.

**Step 10** Click the **Apply Changes** to commit the configuration changes for that feature and distribute the changes through CFS, or click **Undo Changes** to discard the changes for that feature.



**Note** When using CFS with features like DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

To configure the master or seed switch for distribution for each feature using Fabric Manager, follow these steps:

**Step 1** Choose the feature that you want to select the merge master for CFS. For example, choose **Switches > Events > CallHome** from the Physical Attributes pane. The Information pane shows that feature, with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

**Step 2** Check the Master column checkbox for the switch that you want to act as the merge master for this feature.

**Step 3** Click the **Apply Changes** to select this switch as master for future CFS distributions, or click **Undo changes** to discard your unsaved changes.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## A CFS Example Using Device Manager

This procedure is an example of what you see when you use Device Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature's documentation.

To configure feature using Device Manager that uses CFS, follow these steps:

- 
- Step 1** Open the dialog box for any CFS-capable feature. Device Manager checks to see whether CFS is enabled. It also checks to see if there is a lock on the feature by checking for at least one entry in the Owner table.
- Step 2** If CFS is enabled and there is a lock, Device Manager sets the status to “pending” for that feature. You see a dialog box displaying the lock information and you are prompted to continue or cancel the procedure. If you continue, Device Manager remembers the CFS status.
- Step 3** Select **Admin > CFS (Cisco Fabric Services)** to view the username of the CFS lock holder. Click the locked feature and click **Details**.
- Step 4** Click the **Owners** tab and look in the UserName column.



**Note** Device Manager does not monitor the status of the feature across the fabric until the user clicks **Refresh**. If a user on another CFS-enabled switch attempts to configure the same feature, they will not see the “pending” status. However, their configuration changes will be rejected by your switch.

---

- Step 5** If CFS is enabled and there is no lock, Device Manager sets the status to “running” for that feature. You then see a dialog box for the feature. As soon as you perform a creation, deletion, or modification, Device Manager changes the status to “pending” and displays the updated information from the pending database.
- Step 6** View the CFS table for a feature. Device Manager only changes the status to “running” when **commit**, **clear**, or **abort** is selected and applied. Device Manager will not change the status to “pending” if **enable** is selected, because the “pending” status does not apply until the first actual change is made.

The Last Command and Result fields are blank if the last command is noOp.

---



**Note** When using CFS with features like DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## VSAN Configuration

---

You can achieve higher security and greater stability in Fibre Channel fabrics by using Virtual SANs (VSANs). VSANs provide isolation among devices that are physically connected to the same fabric. With VSANs you can create multiple logical SANs over a common physical infrastructure. Each VSAN can contain up to 239 switches and has an independent address space that allows identical Fibre Channel IDs (FC IDs) to be used simultaneously in different VSANs. VSAN members can join the VSAN statically or dynamically.

This chapter includes the following sections:

- [About VSANs, page 13-1](#)
- [Configuring a VSAN, page 13-2](#)
- [Deleting VSANs, page 13-3](#)

### About VSANs

A VSAN is a group of hosts or storage devices that communicate with each other using a virtual topology defined on the physical SAN. Using VSANs, you can build a single topology containing switches, links, and one or more VSANs. Each VSAN in this topology has the same behavior and property of a SAN. A VSAN has the following additional features:

- Multiple VSANs can share the same physical topology.
- The same Fibre Channel IDs (FC IDs) can be assigned to a host in another VSAN, thus increasing VSAN scalability.
- Every instance of a VSAN runs all required protocols such as FSPF, domain manager, and zoning.
- Fabric-related configurations in one VSAN do not affect the associated traffic in another VSAN.
- Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs.

### Default and Isolated VSANs

Up to 256 VSANs can be configured in a switch. Of these, one is a default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Default VSANs

The factory settings for switches in the Cisco MDS 9000 Family have only the default VSAN 1 enabled. If you do not need more than one VSAN for a switch, use this default VSAN as the implicit parameter during configuration. If no VSANs are configured, all devices in the fabric are considered part of the default VSAN. By default, all ports are assigned to the default VSAN.



**Note**

VSAN 1 cannot be deleted, but it can be suspended.

## Isolated VSANs

VSAN 4094 is an isolated VSAN. All non-trunking ports are transferred to this VSAN when the VSAN to which they belong is deleted. This avoids an implicit transfer of ports to the default VSAN or to another configured VSAN. All ports in the deleted VSAN are isolated (disabled).



**Note**

When you configure a port in VSAN 4094 or move a port to VSAN 4094, that port is immediately isolated.



**Caution**


Do not use an isolated VSAN to configure ports.

# Configuring a VSAN

To add and configure a VSAN, follow these steps.

- Step 1** From the Fabric Manager, click **All VSANs** in the Logical Domains pane. The Fabric Manager's Information pane displays VSAN attributes for multiple switches.
- From Device Manager, choose **FC > VSAN** or click the **VSAN** icon on the toolbar as shown in [Figure 13-1](#).

**Figure 13-1 Create VSAN Icon**



The VSAN dialog box in the Device Manager displays VSAN general attributes for a single switch.

- Step 2** From Fabric Manager, click the **Create Row** button on the Information pane toolbar. From Device Manager, click **Create** on the VSAN dialog box.
- You see the Create dialog box.
- Step 3** Complete the fields on this dialog box and click **Create** to add the VSAN.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

To see only the VSANs that a user role is restricted to, you must make the user role and role definition the same on each switch in the fabric. Fabric discovery with Fabric Manager server should be performed by the user defined with that role, unless you are running Fabric Manager server on localhost and the fabrics are not constantly monitored.

## Deleting VSANs

When an active VSAN is deleted, all of its attributes are removed from the running configuration.

To delete a VSAN, follow these steps:

- 
- Step 1** From Fabric Manager, select **All VSANs** from the Logical Domains pane. The VSANs in the fabric are listed in the Information pane on Fabric Manager.
- From Device Manager, choose **FC > VSAN** or click the **VSAN** icon on the toolbar. You see the VSAN dialog box on Device Manager.
- Step 2** Select the VSAN that you want to delete by clicking it.
- Step 3** In the Fabric Manager Information pane toolbar, click the **Delete Row** button.
- In the Device Manager VSAN dialog box, click the **Delete** button.
- You see a confirmation dialog box.
- Step 4** Click **Yes** to confirm the deletion, or **No** to close the dialog box without deleting the VSAN.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Dynamic VSAN Configuration

---

Port VSAN membership on the switch is assigned on a port-by-port basis. By default each port belongs to the default VSAN.

As of Cisco MDS SAN-OS Release 2.x, you can dynamically assign VSAN membership to ports by assigning VSANs based on the device WWN. This method is referred to as the Dynamic Port VSAN Membership (DPVM) feature. DPVM offers flexibility and eliminates the need to reconfigure the VSAN to maintain fabric topology when a host or storage device connection is moved between two Cisco MDS switches. It retains the configured VSAN regardless of where a device is connected or moved.

To assign VSANs statically, see [Chapter 13, “VSAN Configuration.”](#)

This chapter includes the following sections:

- [About DPVM, page 14-1](#)
- [Using the DPVM Setup Wizard, page 14-4](#)
- [Modifying the DPVM Database, page 14-4](#)

### About DPVM

DPVM configurations are based on the port world wide name (pWWN). The pWWN identifies the host or device. A DPVM database contains mapping information for each device pWWN assignment and the corresponding VSAN. Cisco MDS SAN-OS checks the database during a device FLOGI and obtains the required VSAN details.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## DPVM Requirements

Verify the following requirements when using DPVM:

- The interface through which the dynamic device connects to the Cisco MDS switch must be configured as an F port.
- The static port VSAN of the F port should be valid (not isolated, not suspended, and in existence).
- The dynamic VSAN configured for the device in the DPVM database should be valid (not isolated, not suspended, and in existence).



### Note

The DPVM feature overrides any existing static port VSAN membership configuration. If the VSAN corresponding to the dynamic port is deleted or suspended, the port is shut down.



### Note

If you copy the DPVM database and fabric distribution is enabled, you must commit the changes.

To begin configuring the DPVM feature, you must explicitly enable DPVM on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

## DPVM Databases

The DPVM database consists of a series of device mapping entries. Each entry consists of a device pWWN/nWWN assignment along with the dynamic VSAN to be assigned. You can configure a maximum of 16,000 DPVM entries in the DPVM database. This database is global to the switch and fabric and is not maintained for each VSAN.

The DPVM feature uses two databases to accept and implement configurations.

- Configuration (config) database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric.

Changes to the config database are not reflected in the active database until you activate the pending database. This database structure allows you to create multiple entries, review changes, and then activate the pending database.

## DPVM Database Distribution

If the DPVM database is available on all switches in the fabric, devices can be moved anywhere and offer the greatest flexibility. To enable database distribution to the neighboring switches, the database should be consistently administered and distributed across all switches in the fabric. The Cisco MDS SAN-OS software uses the Cisco Fabric Services (CFS) infrastructure to achieve this requirement.

Using the CFS infrastructure, each DPVM server learns the DPVM database from each of its neighboring switches during the ISL bring-up process. Conflicts detected during the local and remote database merge may cause the ISL to be shut down. If you change the database locally, the DPVM server notifies its neighboring switches, and that database is updated by all switches in the fabric.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Config Database Activation

When you explicitly activate the database, the config database becomes the active database. Activation may fail if conflicting entries are found between the config and the currently active database. However, you can force activation to override conflicting entries.

## Copying the DPVM Database

The following circumstances may require that you copy the active database to the config database:

- If the learned entries are only added to the active database.
- If the config database or entries in the config database are accidentally deleted.

## Autolearn Entries

The DPVM database can be configured to automatically learn (autolearn) about new devices within each VSAN. The autolearn feature can be enabled or disabled at any time. Learned entries are created by populating device pWWNs and VSANs in the DPVM active database. The DPVM active database should already be available to enable the autolearn feature.

You can delete any learned entry from the DPVM active database. These entries only become permanent in the active database when you disable the autolearn feature.

The following conditions apply to learned entries:

- If a device logs out while autolearn is enabled, that entry is automatically deleted from the DPVM active database.
- If the same device logs into the switch multiple times through different ports, then the VSAN corresponding to last login is remembered.
- Learned entries do not override previously configured and activated entries.
- Learning is a two-part process:
  - Learning currently logged-in devices—occurs from the time learning is enabled.
  - Learning new device logins— occurs as and when new devices log in to the switch.


*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Using the DPVM Setup Wizard

To use the DPVM Setup Wizard in Fabric Manager to set up dynamic port VSAN membership, follow these steps:

- Step 1** Click the **DPVM Setup Wizard** icon, as shown in [Figure 14-1](#) in the Fabric Manager toolbar.

**Figure 14-1 DPVM Wizard Icon**



You see the Select Master Switch page.

- Step 2** Click the switch you want to be the master switch. This switch controls the distribution of the DPVM database to other switches in the fabric.
- Step 3** Click **Next**. You see the AutoLearn Current End Devices page.
- Step 4** Optionally, click the **Create Configuration From Currently Logged In End Devices** check box if you want to turn on autolearning.
- Step 5** Click **Next**. You see the Edit and Activate Configuration page.
- Step 6** Verify the current or autolearned configuration. Optionally, click **Insert** to add more entries into the DPVM config database.
- Step 7** Click **Finish** to update the DPVM config database, distribute the changes using CFS, and activate the database, or click **Cancel** to exit the DPVM Setup Wizard without saving changes.

## Modifying the DPVM Database

You can modify the DPVM database created with the DPVM Setup wizard. This includes adding or modifying entries, activating the config database, turning autolearning on or off, and clearing unsaved autolearned entries.



### Note

Most tabs in the Information pane for features using CFS are dimmed until you visit the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is visited, the other tabs are activated.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Using the DPVM tables

You can modify the DPVM database through the DPVM Setup Wizard or directly using the DPVM tables in Fabric Manager.

To modify DPVM using the DPVM tables, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>All VSANs &gt; DPVM</b> from the Logical Attributes pane. You see the DPVM tables in the Information pane. |
| <b>Step 2</b> | Choose <b>Config Database</b> to modify existing VSAN entries or choose <b>Create Row</b> to insert a new entry.     |
| <b>Step 3</b> | Choose <b>Active Database</b> or <b>Database Differences</b> to analyze the current DPVM database.                   |
| <b>Step 4</b> | Choose <b>Actions</b> to activate the database or configure autolearning.  |
| <b>Step 5</b> | Choose <b>CFS</b> to select the master switch, and then select <b>commit</b> or <b>abort</b> to discard changes.     |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Zone Configuration

---

Zoning enables you to set up access control between storage devices or user groups. If you have administrator privileges in your fabric, you can create zones to increase network security and to prevent data loss or corruption. Zoning is enforced by examining the source-destination ID field.

As of Cisco MDS SAN-OS Release 2.x, advanced zoning capabilities specified in the FC-GS-4 and FC-SW-3 standards are provided by the Cisco SAN-OS software. You have the option of using the existing basic zoning capabilities or using the advanced, standards-compliant zoning capabilities.

This chapter includes the following sections:

- [Zoning Features, page 15-1](#)
- [Using the Zone Configuration Tool, page 15-3](#)
- [Adding Zone Members, page 15-5](#)
- [Alias Configuration, page 15-6](#)
- [Zone Set Creation, page 15-8](#)
- [Performing Zone Merge Analysis, page 15-15](#)
- [Zone-Based Traffic Priority, page 15-18](#)
- [About LUN Zoning, page 15-19](#)
- [About Read-Only Zones, page 15-21](#)

## Zoning Features

For Cisco MDS SAN-OS Release 2.0(1b), Fabric Manager has added the following to its zoning capabilities:

- Aliases are treated as groups.
- You can have many different types of aliases.
- You can rename zone sets, zones, and aliases.
- You can backup and restore zone database.
- There are enhanced zoning capabilities.

A zone set consists of one or more zones. A zone can be a member of more than one zone set and consists of multiple zone members. Members in a zone can access each other; members in different zones cannot access each other. Devices can belong to more than one zone.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

A zone set can be activated or deactivated as a single entity across all switches in the fabric. Only one zone set can be activated at any time. If zoning is not activated, all devices are members of the default zone. If zoning is activated, any device that is not in an active zone (a zone that is part of an active zone set) is a member of the default zone.

Zoning can be administered from any switch in the fabric. When you activate a zone (from any switch), all switches in the fabric receive the active zone set. Additionally, full zone sets are distributed to all switches in the fabric, if this feature is enabled in the source switch.

If a new switch is added to an existing fabric, zone sets are acquired by the new switch.

Zone changes can be configured nondisruptively. New zones and zone sets can be activated without interrupting traffic on unaffected ports or devices.

Default zone membership includes all ports or WWNs that do not have a specific membership association. Access between default zone members is controlled by the default zone policy.

## Zone Implementation

All switches in the Cisco MDS 9000 Family automatically support the following basic zone features (no additional configuration is required):

- Zones are contained in a VSAN.
- Only active zone sets are distributed.
- Unzoned devices cannot access each other.
- A zone or zone set with the same name can exist in each VSAN.
- Each VSAN has a full database and an active database.
- Active zone sets cannot be changed, without activating a full zone database.
- Active zone sets are preserved across switch reboots.
- Changes to the full database must be explicitly saved.
- Zone reactivation (a zone set is active and you activate another zone set) does not disrupt existing traffic.

If required, you can additionally configure the following zone features:

- Propagate full zone sets to all switches on a per VSAN basis.
- Change the default policy for unzoned members.
- Interoperate with other vendors by configuring a VSAN in the interop mode. You can also configure one VSAN in the interop mode and another VSAN in the basic mode in the same switch without disrupting each other
- Bring E ports out of isolation.

## Zone Configuration

A zone can be configured using one of the following types to assign members:

- pWWN—The WWN of the N or NL port in hex format (for example, 10:00:00:23:45:67:89:ab).
- Fabric port WWN—The WWN of the fabric port name in hex format (for example, 10:00:00:23:45:67:89:ab).
- FC ID—The N port ID in 0xhhhhhh format (for example, 0xce00d1).



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***


- **FC alias**—The alias name is in alphabetic characters (for example, Payroll) and denotes a port ID or WWN. The alias can also include multiple members.
- **Domain ID**—The domain ID is an integer from 1 to 239. A mandatory port number of a non-Cisco switch is required to complete this membership configuration.
- **IP address**—The IP address of an attached device in 32 bytes in dotted decimal format along with an optional subnet mask. If a mask is specified, any device within the subnet becomes a member of the specified zone.
- **Interface**—Interface-based zoning is similar to port-based zoning because the switch interface is used to configure the zone. You can specify a switch interface as a zone member for both local and remote switches. To specify a remote switch, enter the remote switch WWN (sWWN) or the domain ID in the particular VSAN.

## Using the Zone Configuration Tool

To configure zones, read-only zones, and IVR zones using the Zone configuration tool, follow these steps:

- Step 1** From the Fabric Manager toolbar, click the **Zone** icon as shown in [Figure 15-1](#).

**Figure 15-1 Zone Icon**



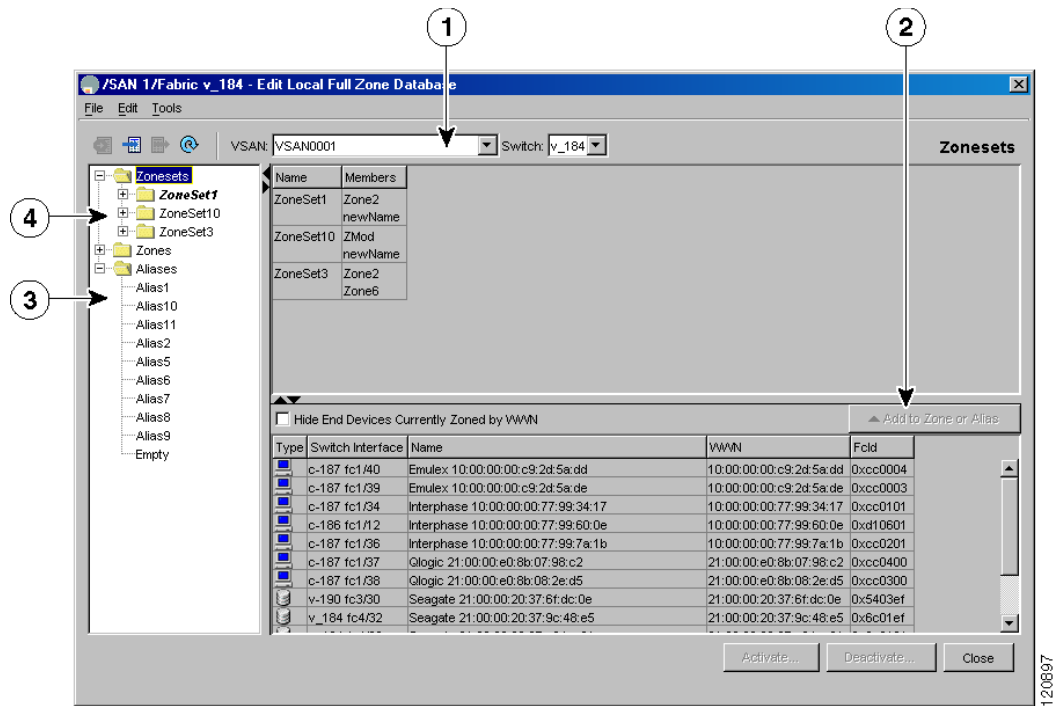
- Step 2** Select the VSAN where you want to configure zone sets, zones, or add members to a zone.
- Step 3** Click **Zoneset** and click the **Create Row** icon to make a new zone set.
- Step 4** Click **Zones** and click the **Create Row** icon to make a new zone.
- Optionally, check **Read Only** check box if you want the zone to permit reads and deny writes.
  - Optionally, check the **Permit QoS traffic with Priority** check box and set the QoS priority from the drop-down menu.
  - Optionally, check the **Restrict Broadcast frames to Zone Members** check box.
- Step 5** Click **Aliases** and click the **Create Row** icon to create a new device alias.
- Step 6** Click a zone and click the **Create Row** icon to create a new zone member.
- Select the zone member type (for example, FC ID, pWWN) at set the appropriate name or ID.
  - Click **Add** to add the zone member to the zone or click **Cancel** to close the dialog box without adding a new zone member.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Edit Full Zone Database Overview

For version 2.0, there are interface changes to the Edit Full Zone Database screen, which is shown in Figure 15-2.

**Figure 15-2 Edit Full Zone Database Screen**



1	You can display information by VSAN by using the pull-down menu without having to get out of the screen, selecting a VSAN, and re-entering.	3	You can add zoning characteristics based on alias in different folders.
2	You can use the <b>Add to zone or alias</b> button to move devices up or down by alias or by zone.	4	You can triple-click to rename zone sets, zones, or aliases in the tree.

## Zone Database Information

If required, you can clear configured information stored in the zone server database.



### Note

Clearing a zone set only erases the full zone database, not the active zone database.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring a Zone

**Note**

Interface-based zoning only works with Cisco MDS 9000 Family switches. Interface-based zoning does not work if **interop** mode is configured in that VSAN.

**Tip**

If you do not provide an sWWN, the software automatically uses the local sWWN.

Zones are configured within VSANs, but you can configure zones without configuring any VSANs by configuring them within the default VSAN. The Logical tab displays the VSANs configured in the currently discovered fabric. Note that zone information must always be identical for all the switches in the network fabric.

To create zones, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.
- If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected in the Logical Attributes pane, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.
- You see the Edit Local Full Zone Database dialog box for the VSAN you selected.
- Step 2** Right-click the **Zone** folder in the Edit Local Full Zone Database dialog box for that VSAN and select **Insert** to add a zone.
- You can specify that the zone be a read-only zone by checking the **Read Only** check box. (See the [“About Read-Only Zones”](#) section on page 15-21.)
- 

## Viewing Zone Statistics

To monitor zone statistics from the Zone Server, choose **VSANxxx > Domain Manager** from the Fabric Manager menu tree. The zone information is displayed in the Information pane. Click the **Statistics** tab to see the statistics information for the switches in the zone.

## Adding Zone Members

Once you have created a zone, you can add members to the zone. You can add members using multiple port identification types. See the [“Zone Configuration”](#) section on page 15-2.

To add a member to a zone, follow these steps:

- 
- Step 1** Click the **Zones** folder from the Logical Attributes pane. Right-click the folder for the zone to which you want to add members, and then choose **Insert** from the pop-up menu.
- You see the Add Member to Zone dialog box.
- Step 2** Click the check box to the left of the port identification type you want to add.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 3** Select one of the port identifier options in the dialog box and click **Add** to add it to the zone.  
You see the member in the zone server database in the lower frame.

**Step 4** Repeat these steps to add other members to the zone.



**Note** When configuring a zone member, you can specify that a single LUN has multiple IDs depending on the operating system. You can select from six different operating systems.

## Displaying Zone Membership Information

To display zone membership information for members assigned to zones, follow these steps:

**Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.  
If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

**Step 2** Click the **Zones** folder. The right pane lists the members for each zone.



**Note** The default zone members are explicitly listed only when the default zone policy is configured as permit. When the default zone policy is configured as deny, the members of this zone are not shown. See the [“Viewing Zone Statistics” section on page 15-5](#).

## Alias Configuration

You can assign an alias name and configure an alias member using either the FC ID, fabric port WWN (fWWN), pWWN, domain ID and port number, interface ID, IP address, or symbolic node name values.



**Tip** Cisco MDS SAN-OS Release 1.3(4) or later supports a maximum of 2048 aliases per VSAN.

## Creating Zones with Aliases

To create a zone with aliases, follow these steps:

**Step 1** Select **Edit Local Full Zone Database...** from the Zone menu.

You see the Select VSAN dialog box.

**Step 2** Select the VSAN on which you want to create the zone and click **OK**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You see the zone information for that VSAN.

**Step 3** Right-click the **Aliases** folder in the left window pane and select **Insert**.

You see the Create Alias dialog box.

**Step 4** Enter the alias name and click **OK** to create the alias.

**Step 5** Right-click the newly created alias and select **Insert**. You can add/associate multiple pWWNs and fWWNs to the same alias name. The pWWNs do not have to be attached to the fabric you are currently managing.

**Step 6** Click **Add** to add this entity to this alias.

**Step 7** Right-click the **Zones** folder in the left pane and select **Insert**.

**Step 8** Name the zone as desired and click **OK**.

**Step 9** Select the newly created zone from the right pane and select **Insert**. You see the Add Member Dialog box.

**Step 10** Select the **Alias** radio button. Type the name of the alias you want to associate with this zone or select the ... button to see a list of aliases to select from. Click **OK**.

**Step 11** Add the zone to a zone set and activate it accordingly.

---

## Viewing Aliases

Aliases are assigned per port.

To view zone aliases, follow these steps:

---

**Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.

If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.

You see the Edit Local Full Zone Database window for the VSAN you selected.

**Step 2** Click the **Zones** folder for the zone you are interested in. The aliases for that zone are listed in the right pane.


---

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Zone Set Creation

In [Figure 15-3](#), two separate sets are created, each with its own membership hierarchy and zone members.

**Figure 15-3 Hierarchy of Zone Sets, Zones, and Zone Members**



Zones provide a mechanism for specifying access control, while zone sets are a grouping of zones to enforce access control in the fabric. Either zone set A or zone set B can be activated (but not together).



### Note

You can specify multiple zone members on multiple lines at the switch prompt.



### Tip

Zone sets are configured with the names of the member zones. If the zone set is in a configured VSAN, you must also specify the VSAN.

## Active and Full Zone Set Considerations

Before configuring a zone set, consider the following guidelines:

- Each VSAN can have multiple zone sets but only one zone set can be active at any given time.
- When you create a zone set, that zone set becomes a part of the full zone set.
- When you activate a zone set, a copy of the zone set from the full zone set is used to enforce zoning, and is called the active zone set. An active zone set cannot be modified. A zone that is part of an active zone set is called an active zone.
- The administrator can modify the full zone set even if a zone set with the same name is active.
- When the activation is done, the active zone set is automatically stored in persistent configuration. This enables the switch to preserve the active zone set information across switch resets.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- All other switches in the fabric receive the active zone set so they can enforce zoning in their respective switches.
- Hard and soft zoning are implemented using the active zone set. Modifications take effect during zone set activation.
- An FC ID or Nx port that is not part of the active zone set belongs to the default zone and the default zone information is not distributed to other switches.

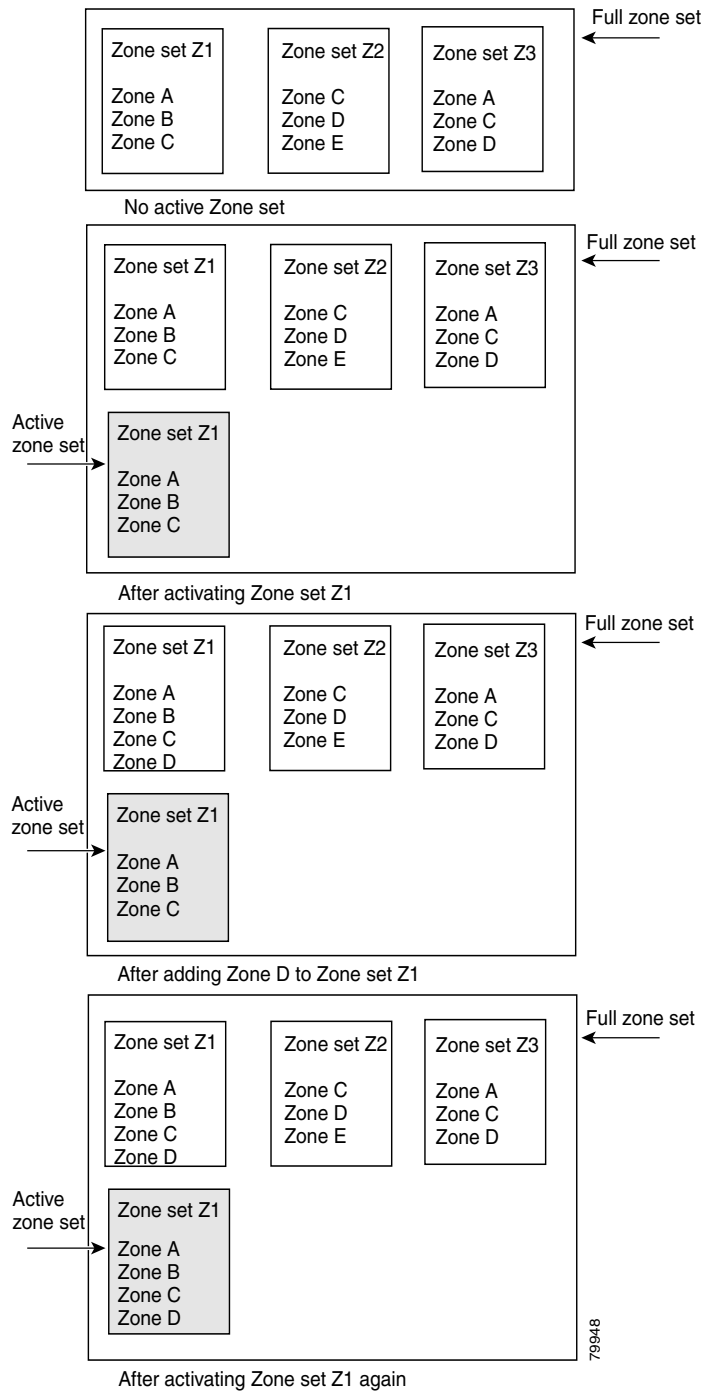
**Note**

If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated. You do not need to explicitly deactivate the currently active zone set before activating a new zone set.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

Figure 15-4 shows a zone being added to an active zone set.

**Figure 15-4 Active and Full Zone Sets**





***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Creating Zone Sets

To create zone sets, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.
- If you chose **Zone > Edit Local Full Zone Database**, or the **All VSANs** folder is selected, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.
- You see the Edit Local Full Zone Database window for the VSAN you selected.
- Step 2** Right-click the **Zonesets** folder in the Edit Local Full Zone Database dialog box for that VSAN and select **Insert** to add a zone set.
- You can activate the zone set after creation by clicking the **Activate** button. This button appears when you right-click the newly created zone set. This configuration is distributed to the other switches in the network fabric.



---

**Note** When you confirm the activate operation, the current running configuration is saved to the startup configuration. This permanently saves any changes made to the running configuration (not just zoning changes).

---

## Adding Zones to a Zone Set

To add a zone to a zone set from the Edit Local Full Zone Database window, drag and drop the zone to the folder for the zone set.

Alternatively, follow these steps:

- 
- Step 1** Click the **Zone sets** folder and then right-click the folder for the zone set to which you want to add a zone and choose **Insert** from the pop-up menu.
- You see the Zone dialog box. You can filter the entries in the Zone dialog box by entering the first few letters of the zones you are searching for in the top text box in the Zone dialog box.
- Step 2** Select the zone that you want to add to the zone set and click **Add**.
- The zone is added to the zone set in the zone database.
- 

## Activating Zone Sets

Once zones and zone sets have been created and populated with members, you must activate the zone set. Note that only one zone set can be activated at any time. If zoning is activated, any member that is not assigned to an active zone belongs to the default zone. If zoning is not activated, all members belong to the default zone.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To activate a zone set, follow these steps:

---

**Step 1** Right-click the zone set in the Edit Local Full Database dialog box.

**Step 2** Click **Activate**.

You see the zone set in the Active Zone Set folder.




---

**Note** If one zone set is active and you activate another zone set, the currently active zone set is automatically deactivated.

---

## Deactivating Zone Sets

To deactivate a zone set, follow these steps:

---

**Step 1** Right-click the zone set in the Edit Full Database dialog box.

**Step 2** Click **Deactivate**.

You see the zone set removed from the Active Zone Set folder.

---

## Creating Additional Zones and Zone Sets

To create additional zones and zone sets, follow these steps:

---

**Step 1** With the Edit Full Database dialog box open, right-click the **Zones** folder and choose **Insert** from the pop-up menu.

**Step 2** Enter the zone name in the dialog box that appears and click **OK** to add the zone.

The zone is automatically added to the zone database.

**Step 3** Right-click the **Zonesets** folder in the Edit Full Database dialog box, and choose **Insert**.

**Step 4** Enter the zone set name in the dialog box that appears and click **OK** to add the zone set.

The zone set is automatically added to the zone database.

---

## Cloning Zones and Zone Sets

You can make a copy and then edit it without altering the existing active zone set. You can copy an active zone set from the bootflash: directory, volatile: directory, or slot0, to one of the following areas:

- To the full zone set
- To a remote location (using FTP, SCP, SFTP, or TFTP).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The active zone set is not part of the full zone set. You cannot make changes to an existing zone set and activate it, if the full zone set is lost or is not propagated.

To clone a zone or zone set from the Edit Local Full Zone Database window, follow these steps:

- 
- Step 1** Select the **Zones** or **Zonesets** folder, right-click the folder for the zone or zone set that you want to clone, and choose **Clone** from the pop-up menu.
- Step 2** Enter the name of the cloned zone or zone set.
- By default, the dialog box displays the selected zone name by prepending the original zone name with "Cloned" (for example, ClonedZone1) and selects the read-only zone state to match the cloned zone.
- Step 3** Click **OK** to add the cloned zone to the zone database.
- 

**Caution**

Copying an active zone set to a full zone set may overwrite a zone with the same name, if it already exists in the full zone set database.

---

## Deleting Zones, Zone Sets, and Aliases

To delete zones, zone sets, or aliases, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.
- If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.
- You see the Edit Local Full Zone Database window for the VSAN you selected.
- Step 2** Select the zone, zone set, or alias you want to delete.
- Step 3** Right-click the object and choose **Delete** from the pop-up menu, or click the **Delete** button.
- The selected object is deleted from the zone database.
- 

## Zone Enforcement

Zoning can be enforced in two ways: soft and hard. Each end device (N port or NL port) discovers other devices in the fabric by querying the name server. When a device logs in to the name server, the name server returns the list of other devices that can be accessed by the querying device. If an Nx port does not know about the FC IDs of other devices outside its zone, it cannot access those devices.

In soft zoning, zoning restrictions are applied only during interaction between the name server and the end device. If an end device somehow knows the FC ID of a device outside its zone, it can access that device.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Hard zoning is enforced by the hardware on each frame sent by an Nx port. As frames enter the switch, source-destination IDs are compared with permitted combinations to allow the frame at wirespeed. Hard zoning is applied to all forms of zoning.



**Note**

Hard zoning enforces zoning restrictions on every frame, and prevents unauthorized access.

Switches in the Cisco MDS 9000 Family support both hard and soft zoning.

## The Default Zone

Each member of a fabric (in effect a device attached to an Nx port) can belong to any zone. If a member is not part of any active zone, it is considered to be part of the default zone. Therefore, if no zone set is active in the fabric, all devices are considered to be in the default zone. Even though a member can belong to multiple zones, a member that is part of the default zone cannot be part of any other zone. The switch determines whether a port is a member of the default zone when the attached port comes up.



**Note**

Unlike configured zones, default zone information is not distributed to the other switches in the fabric.

Traffic can either be permitted or denied among members of the default zone. This information is not distributed to all switches; it must be configured in each switch.



**Note**

When the switch is initialized for the first time, no zones are configured and all members are considered to be part of the default zone. Members are not permitted to talk to each other.

Configure the default zone policy on each switch in the fabric. If you change the default zone policy on one switch in a fabric, be sure to change it on all the other switches in the fabric.



**Note**

The default settings for default zone configurations can be changed.

## Configuring the Default Zone Policy

The default zone members are explicitly listed when the default policy is configured as **permit** or when a zone set is active. When the default policy is configured as **deny**, the members of this zone are not explicitly enumerated when you deactivate the zone set.

You can change the default zone policy for any VSAN by choosing **VSANxxx > Default Zone** from the Fabric Manager menu tree and clicking the **Policies** tab. It is recommended that you establish connectivity among devices by assigning them to a non-default zone.

The active zone set is shown in italic type. After you have made changes to the active zone set and before you activate the changes, the zone set is shown in boldface italic type.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To permit or deny traffic to members in the default zone from the Zone Server, follow these steps:

- 
- Step 1** Choose **VSANxxx > Default Zone** from the Fabric Manager Logical Domains menu tree, and click the **Policies** tab in the Information pane.
- You see the zone information in the Information pane.
- Step 2** Click the **Default Zone Behavior** field and choose either **permit** or **deny** from the pull-down menu.
- 

## Performing Zone Merge Analysis

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. You can perform a zone merge analysis prior to merging the switches to see if the merge will succeed or fail.

To perform a zone merge analysis, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Merge Analysis** from the Zone menu.
- You see the Zone Merge Analysis dialog box.
- Step 2** Select the first switch to be analyzed from the Check Switch 1 drop-down list.
- Step 3** Select the second switch to be analyzed from the And Switch 2 drop-down list.
- Step 4** Enter the VSAN ID where the zone set merge failure occurred in the For Active Zoneset Merge Problems in VSAN Id field.
- Step 5** Click **Analyze** to analyze the zone merge. Click **Clear** to clear the analysis data from the Zone Merge Analysis dialog box.
- 

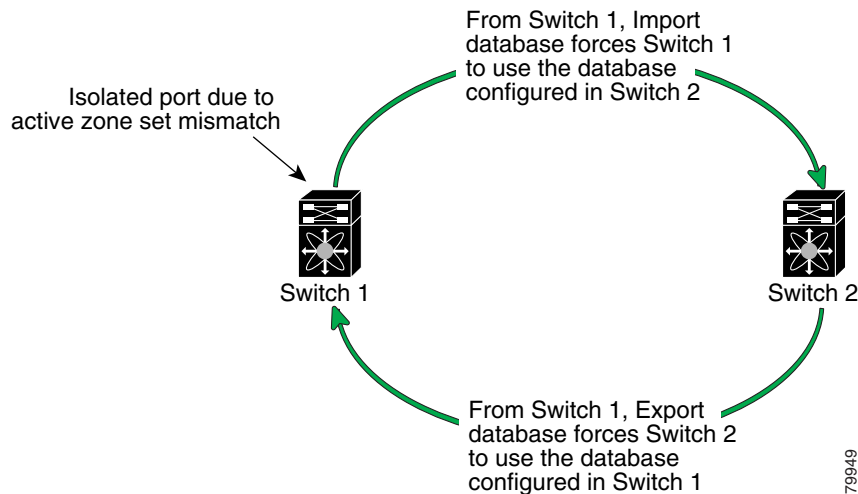
## Recovering from Link Isolation

When two switches in a fabric are merged using a TE or E port, these TE and E ports may become isolated when the active zone set databases are different between the two switches or fabrics. When a TE port or an E port become isolated, you can recover that port from its isolated state using one of three options:

- Import the neighboring switch's active zone set database and replace the current active zone set (see [Figure 15-5](#)).
- Export the current database to the neighboring switch (see [Figure 15-5](#)).
- Manually resolve the conflict by editing the full zone set, activating the corrected zone set, and then bringing up the link.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 15-5 Importing and Exporting the Database**



## Importing Zone Sets



### Note

Importing from one switch and exporting from another switch can lead to isolation again.

You can import active zone sets (do a Merge Fail Recovery) if the cause of an ISL failure is a zone merge failure.

To import an active zone set, follow these steps:

- Step 1** From Fabric Manager, choose **Zone > Merge Fail Recovery** from the Zone menu.  
You see the Zone Merge Failure Recovery dialog box.
- Step 2** Select the **Import Active Zoneset** radio button.
- Step 3** Select the switch from which to import the zone set information from the drop-down list.
- Step 4** Select the VSAN from which to import the zone set information from the drop-down list.
- Step 5** Select the interface to use for the import process.
- Step 6** Click **OK** to import the active zone set, or click **Close** to close the dialog box without importing the active zone set.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Exporting Active Zone Sets

You can export active zone sets (do a Merge Fail Recovery) if the cause of an ISL failure is a zone merge fail.

To export an active zone set, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Merge Fail Recovery** from the Zone menu.  
You see the Zone Merge Failure Recovery dialog box.
  - Step 2** Select the **Export Active Zoneset** radio button.
  - Step 3** Select the switch to which to export the zone set information from the drop-down list.
  - Step 4** Select the VSAN to which to export the zone set information from the drop-down list.
  - Step 5** Select the interface to use for the export process.
  - Step 6** Click **OK** to export the active zone set, or click **Close** to close the dialog box without exporting the active zone set.
- 

## Full Zone Set Propagation

All switches in the Cisco MDS 9000 Family distribute active zone sets when new E port links come up or when a new zone set is activated in a VSAN. The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

To propagate the full zone set from Fabric Manager, follow these steps:

- 
- Step 1** Select **VSANxxx > ZoneSetxx** from the Logical Domains pane. You see the zone set configuration in the Information pane.
  - Step 2** Select the **Policies** tab.
  - Step 3** Set the propagation column to **fullZoneset** from the drop-down menu.
  - Step 4** Click **Apply Changes** to propagate the full zone set, or click **Undo Changes** to discard any changes you made.
- 

## One-Time Distribution

To propagate a one-time distribution of the full zone set from Fabric Manager, follow these steps:

- 
- Step 1** Select **Zone > Edit Local Full Zone Database** from the main menu.
  - Step 2** Select the appropriate VSAN from the list. You see the Edit Local Full Zone Set configuration tool.
  - Step 3** Click **Distribute** to distribute the full zone set across the fabric.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Copying a Full Zone Database

You can recover a database by copying the active zone database or the full zone database.

To copy a zone database, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Copy Full Zone Database** from the Zone menu.  
You see the Recover Full Zone Database dialog box.
  - Step 2** Select the **Copy Active** or the **Copy Full** radio button, depending on which type of database you want to copy.
  - Step 3** Select the source VSAN from which to copy the information from the drop-down list.
  - Step 4** If you selected **Copy Full**, select the source switch and the destination VSAN from those drop-down lists.
  - Step 5** Select the destination switch from the drop-down list.
  - Step 6** Click **Copy** to copy the database, or click **Close** to close the dialog box without copying.
- 

## Migrating a Non-MDS Database

To use the Zone Migration Wizard to migrate a non-MDS database, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Migrate Non-MDS Database** from the Zone menu.  
You see the Zone Migration Wizard.
  - Step 2** Follow the prompts in the wizard to migrate the database.
- 

## Zone-Based Traffic Priority

As of Cisco MDS SAN-OS 2.0, the zoning feature provides an additional segregation mechanism to prioritize select zones in a fabric and set up access control between devices. Using this feature, you can configure the Quality of Service (QoS) priority as a zone attribute. You can assign the QoS traffic priority attribute to be **high**, **medium**, or **low**. By default, zones with no specified priority are implicitly assigned a **low** priority.

To use this feature, you need to obtain the ENTERPRISE\_PKG license (see [Chapter 9, “Obtaining and Installing Licenses”](#)) and you must enable QoS in the switch.

This feature allows SAN administrators to configure QoS in terms of a familiar data flow identification paradigm. You can configure this attribute on a zone-wide basis, rather than between zone members.

## Configuring Zone QoS and Broadcast Attributes

QoS attribute-specific configuration changes take effect when you activate the zone set of the associated zone.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

If a member is part of two zones with two different QoS attributes, the higher QoS value is implemented. This situation does not arise in the VSAN-based QoS as the first matching entry is implemented.

As of Cisco MDS SAN-OS Release 2.0(1b), you can configure broadcast frames in the basic zoning mode. By default, broadcast zoning is disabled. When enabled, broadcast frames are sent to all Nx Ports. Broadcast zoning can only be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 2.0(1b) or later.

**Tip**

If any NL port attached to an FL port shares a broadcast zone with the source of the broadcast frame, then the frames are broadcast to all devices in the loop.

**Caution**

If Broadcast zoning is implemented in a switch, you cannot configure the interop mode in that VSAN.

To configure the zone QoS or broadcast attributes in Fabric Manager, follow these steps:

- Step 1** Choose **VSANxxx > <zone set name>** from the Fabric Manager Logical Domains menu tree, and click the **Policies** tab in the Information pane.  
You see the Zone policy information in the Information pane.
- Step 2** Check the **QoS** check box to enable QoS on the default zone.
- Step 3** Click **QoS Priority** and choose **low**, **medium**, or **high** from the pull-down menu.
- Step 4** Check the **Broadcast** check box to enable broadcast frames on the default zone.
- Step 5** Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to discard these changes.

## About LUN Zoning

Logical unit number (LUN) zoning is a feature specific to switches in the Cisco MDS 9000 Family.

**Caution**

LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure the interop mode in that switch.

**Note**

LUN zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or earlier.

A storage device can have multiple LUNs behind it. If the device port is part of a zone, a member of the zone can access any LUN in the device. With LUN zoning, you can restrict access to specific LUNs associated with a device.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Note**

When LUN 0 is not included within a zone, then, as per standards requirements, control traffic to LUN 0 (for example, REPORT\_LUNS, INQUIRY) is supported, but data traffic to LUN 0 (for example, READ, WRITE) is denied.


- Host H1 can access LUN 2 in S1 and LUN 0 in S2. It cannot access any other LUNs in S1 or S2.
- Host H2 can access LUNs 1 and 3 in S1 and only LUN 1 in S2. It cannot access any other LUNs in S1 or S2.

**Note**

Unzoned LUNs automatically become members of the default zone.

Figure 15-6 shows a LUN-based zone example.

**Figure 15-6 LUN Zoning Access**



## Configuring a LUN-Based Zone

To create LUN-based zones, follow these steps:

- Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.  
If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.  
You see the Edit Local Full Zone Database window for the VSAN you selected.
- Step 2** Right-click the **Zones** folder in the Edit Local Full Zone Database dialog for that VSAN and select **Insert** to add a zone.  
You can specify that the zone be a read-only zone by checking the Read Only check box. (For more information on read-only zones, see the [“About Read-Only Zones”](#) section on page 15-21.)
- Step 3** Select either **WWN** or **FCID** radio button for the Zone By options to create a LUN-based zone.
- Step 4** Check the **LUN** check box and add the LUNs for this zone in the text box.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Click **Add** to add this LUN-based zone or **Close** to close the dialog box without adding the LUN-based zone.
- 

## Assigning LUNs to Storage Subsystems

LUN masking and mapping restricts server access to specific LUNs. If LUN masking is enabled on a storage subsystem and if you want to perform additional LUN zoning in a Cisco MDS 9000 Family switch, obtain the LUN number for each Host Bus Adapter (HBA) from the storage subsystem and then configure the LUN-based zone procedure provided earlier.



### Note

Refer to the relevant user manuals to obtain the LUN number for each HBA.



### Caution

If you make any errors when configuring this scenario, you are prone to loose data.

## About Read-Only Zones



### Note

Read-only zoning can be implemented in Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 1.2(x) or above.

By default, an initiator has both read and write access to the target's media when they are members of the same Fibre Channel zone. The read-only zone feature allows members to have only read access to the media within a read-only Fibre Channel zone.

You can also configure LUN zones as read-only zones.

## Guidelines to Configure Read-Only Zones

Any zone can be identified as a read-only zone. By default all zones have read-write permission unless explicitly configured as a read-only zone.

Follow these guidelines when configuring read-only zones:

- If read-only zones are implemented, the switch prevents write access to user data within the zone.
- If two members belong to a read-only zone and to a read-write zone, read-only zone has priority and write access is denied.
- LUN zoning can only be implemented in Cisco MDS 9000 Family switches. If LUN zoning is implemented in a switch, you cannot configure interop mode in that switch.
- Read-only volumes are not supported by some operating system and file system combinations (for example, Windows NT or Windows 2000 and NTFS file system). Volumes within read-only zones are not available to such hosts. However, if these hosts are already booted when the read-only zones are activated, then read-only volumes are available to those hosts.

The read-only zone feature behaves as designed if FAT16 or FAT32 file system is used with the above-mentioned Windows operating systems.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring Read-Only Zones

To create read-only zones, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > Edit Local Full Zone Database** from the Zone menu, or right-click a VSAN folder in the Logical tab and choose **Edit Local Full Zone Database** from the pop-up menu.
- If you chose **Zone > Edit Local Full Zone Database**, then you see the Select VSAN dialog box. Select the VSAN and click **OK**.
- You see the Edit Local Full Zone Database window for the VSAN you selected.
- Step 2** Right-click the **Zones** folder in the Edit Local Full Zone Database dialog box for that VSAN and select **Insert** to add a zone.
- Step 3** Check the **Read Only** check box to create a read-only zone.
- 

## Backing Up and Restoring Zones

You can back up the zone configuration to a workstation using TFTP. This zone backup file can then be used to restore the zone configuration on a switch. Restoring the zone configuration overwrites any existing zone configuration on a switch.

To backup or restore the full zone configuration using Fabric Manager, follow these steps:

- Step 4** Choose **File > Backup** to back up the existing zone configuration to a workstation using TFTP.
- Step 5** Choose **File > Restore** to restore a saved zone configuration. You can optionally edit this configuration before restoring it to the switch.
-



## Inter-VSAN Routing Configuration

---

This chapter explains the inter-VSAN routing (IVR) feature and provides details on sharing resources across VSANs using IVR management interfaces provided in the switch.

This chapter includes the following sections:

- [Inter-VSAN Routing, page 16-1](#)
- [Using the IVR Zone Wizard, page 16-7](#)
- [Modifying IVR, page 16-8](#)
- [Enabling IVR Without NAT, page 16-10](#)
- [IVR Zones and IVR Zone Sets, page 16-13](#)
- [IVR Interoperability, page 16-17](#)

## Inter-VSAN Routing

Virtual SANs (VSANs) improve storage area network (SAN) scalability, availability, and security by allowing multiple Fibre Channel SANs to share a common physical infrastructure of switches and ISLs. These benefits are derived from the separation of Fibre Channel services in each VSAN and isolation of traffic between VSANs. Data traffic isolation between the VSANs also inherently prevents sharing of resources attached to a VSAN, like robotic tape libraries. Using IVR, resources across VSANs are accessed without compromising other VSAN benefits.

This section includes the following topics:

- [Understanding IVR, page 16-1](#)
- [IVR Terminology, page 16-2](#)
- [Fibre Channel Header Modifications, page 16-3](#)
- [IVR NAT, page 16-3](#)
- [IVR VSAN Topology, page 16-4](#)

## Understanding IVR

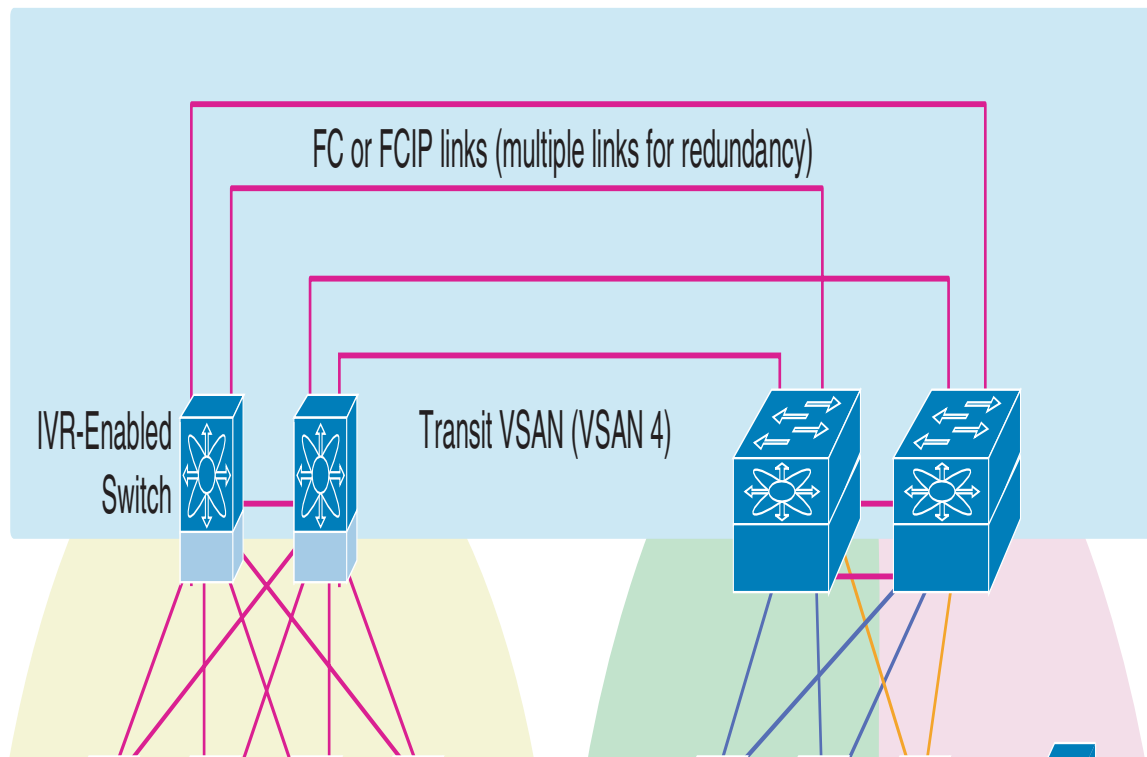
Data traffic is transported between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric. Fibre Channel control traffic does not flow between VSANs, nor can initiators access any resource across VSANs aside from the designated ones. Valuable resources such as tape libraries are easily shared across VSANs without compromise.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

IVR is in compliance with Fibre Channel standards and incorporates third-party switches, however, IVR-enabled VSANs may have to be configured in one of the interop modes.

IVR is not limited to VSANs present on a common switch. Routes that traverse one or more VSANs across multiple switches can be established, if necessary, to establish proper interconnections. IVR used in conjunction with FCIP provides more efficient business continuity or disaster recovery solutions (see [Figure 16-1](#)).

**Figure 16-1** Traffic Continuity Using IVR and FCIP



## IVR Terminology

The following terms are used in this chapter.

- **Native VSAN**—The VSAN to which an end device logs on is the native VSAN for that end device.
- **Inter-VSAN zone (IVR zone)**—A set of end devices that are allowed to communicate across VSANs within their interconnected SAN fabric. This definition is based on their port world wide names (pWWNs) and their native VSAN associations. You can configure up to 2,000 IVR zones and 10,000 IVR zone members in the fabric from any switch in the Cisco MDS 9000 Family.
- **Inter-VSAN zone sets (IVR zone sets)**—One or more IVR zones make up an IVR zone set. You can configure up to 32 IVR zone sets on any switch in the Cisco MDS 9000 Family. Only one IVR zone set can be active at any time.
- **IVR path**—An IVR path is a set of switches and Inter-Switch Links (ISLs) through which a frame from one end-device in one VSAN can reach another end-device in some other VSAN. Multiple paths can exist between two such end-devices.
- **IVR-enabled switch**—A switch in which the IVR feature is enabled.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Edge VSAN—A VSAN that initiates (source edge-VSAN) or terminates (destination edge-VSAN) an IVR path. Edge VSANs may be adjacent to each other or they may be connected by one or more transit VSANs. In [Figure 16-1](#), VSANs 1, 2, and 3 are edge VSANs.



---

**Note** An edge VSAN for one IVR path can be a transit VSAN for another IVR path.

---

- Transit VSAN—A VSAN that exists along an IVR path from the source edge VSAN of that path to the destination edge VSAN of that path. In [Figure 16-1](#), VSAN 4 is a transit VSAN.



---

**Note** When the source and destination edge VSANs are adjacent to each other, then a transit VSAN is not required between them.

---

- Border switch—An IVR-enabled switch that is a member of two or more VSANs. Border switches in [Figure 16-1](#) span two or more different color-coded VSANs.
- Edge switch—A switch to which a member of an IVR zone has logged in. Edge switches are oblivious to the IVR configurations in the border switches. Edge switches need not be IVR enabled.

## Fibre Channel Header Modifications

IVR works by virtualizing the remote end devices in the native VSAN using a virtual domain. When IVR is configured to link end devices in two disparate VSANs, the IVR border switches are responsible for modifying the Fibre Channel headers for all communication between the end devices. The sections of the Fibre Channel frame headers that are modified include:

- VSAN number
- Source FCID
- Destination FCID

When a frame goes from the initiator to the target, the Fibre Channel frame header is modified such that the initiator VSAN number is changed to the target VSAN number. If IVR Network Address Translation (NAT) is enabled, then the source and destination FCIDs are also translated at the edge border switch. If IVR NAT is not enabled, then you must configure unique domain IDs for all switches involved in the IVR path.

## IVR NAT

Cisco MDS SAN-OS Release 2.1(1a) introduces IVR NAT, which allows you to set up IVR in a fabric without needing unique domain IDs on every switch in the IVR path. When IVR NAT is enabled, the virtualized end device that appears in the native VSAN uses a virtual domain ID that is unique to the native VSAN.



**Note**

---

IVR NAT requires Cisco MDS SAN-OS Release 2.1(1a) or later on all switches in the fabric performing IVR. If you have isolated switches with an earlier release that are involved in IVR, you must remove any isolated fabrics from monitoring by Fabric Manager server and then re-open the fabric to use IVR NAT. See [“Removing a Fabric from Monitoring”](#) section on page 2-8.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## IVR VSAN Topology

IVR uses a configured IVR VSAN topology to determine how to route traffic between the initiator and the target across the fabric. You can configure this IVR VSAN topology manually on an IVR-enabled switch and distribute the configuration using CFS in Cisco MDS SAN-OS Release 2.0(1b) or later. Alternately, in Cisco MDS SAN-OS Release 2.1(1a) or later, you can configure IVR topology in auto mode. Prior to Cisco MDS SAN-OS Release 2.0(1b), you need to manually copy the IVR VSAN topology to each switch in the fabric.

Auto mode automatically builds the IVR VSAN topology and maintains the topology database when fabric reconfigurations occur. Auto mode distributes the IVR VSAN topology to IVR-enabled switches using CFS.

Using auto mode, you no longer need to manually update the IVR VSAN topology when reconfigurations occur in your fabric. If a manually configured IVR topology database exists, auto mode initially uses that topology information. This reduces disruption in the network by gradually migrating from the user-specified topology database to the automatically learned topology database. User configured topology entries that are not part of the network are aged out in about three minutes. New entries that are not part of the user configured database are added as they are learned from the network.

**Note**

IVR topology in auto mode requires Cisco MDS SAN-OS Release 2.1(1a) or later and enabling CFS for IVR on all switches in the fabric.

## Autonomous Fabric ID

The autonomous fabric ID (AFID) distinguishes segmented VSANS (that is, two VSANs that are logically and physically separate but have the same VSAN number). Cisco MDS SAN-OS Release 2.1(1a) introduces support for AFIDs from 1 through 64. AFIDs are used in conjunction with auto mode to allow segmented VSANS in the IVR VSAN topology database. You can configure up to 64 AFIDs.

The AFID can be configured individually for each switch and list of VSANs, or the default AFID can be configured for each switch.

**Note**

Two VSANs with the same VSAN number but different AFIDs are counted as two VSANs out of the total 128 VSANs allowed in the fabric.

## Service Groups

Cisco MDS SAN-OS Release 2.1(1a) introduces service groups as a way to limit the control traffic associated with distributing the IVR VSAN topology learned in auto mode. A services group lists AFIDs and the VSANs associated with each AFID. When the IVR configuration is distributed, CFS uses the service group to limit the number of switches to which it sends the new IVR VSAN topology database. Currently, you can configure one service group for the fabric.

**Note**

You must update the service group and distribute it using CFS whenever a fabric reconfiguration affects an IVR-enabled switch.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Using IVR NAT and Auto Topology

Before configuring an IVR SAN fabric to use IVR NAT and auto-topology, consider the following guidelines:

- Configure IVR only in the relevant switches.
- Enable CFS for IVR on all switches in the fabric.
- Verify all switches in the fabric are running Cisco MDS SAN-OS Release 2.1(1a) or later.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package if you have Cisco MDS SAN-OS Release 2.1(1a) or later and one active IPS card for this feature.



### Tip

If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.



### Note

IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

## Transit VSAN Guidelines

Consider the following guidelines for transit VSANs:

- Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also double-up as an edge VSAN in another IVR zone.

## Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 2.1(1a) or later.
- A border switch must be a member of two or more VSANs.
- A border switch that facilitates IVR communications must be IVR enabled.
- IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVR zone members.

The VSAN topology configuration updates automatically when a border switch is added or removed.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Service Group Guidelines

If you use service groups with IVR auto topology, you should enable IVR and configure your service group first, then distribute them with CFS before setting the IVR topology in auto mode.

## Using IVR Without IVR NAT or Auto Topology

Before configuring an IVR SAN fabric without IVR in NAT mode or IVR topology in auto mode, consider the following guidelines:

- Configure unique domain IDs across all VSANs and switches participating in IVR operations if you are not using IVR NAT. The following switches participate in IVR operations:
  - All edge switches in the edge VSANs (source and destination)
  - All switches in transit VSANs
- Configure IVR only in the relevant border switches.
- Acquire a mandatory Enterprise License Package or SAN-EXTENSION license package if you have Cisco MDS SAN-OS Release 2.1(1a) or later and one active IPS card for this feature.



### Tip

If you change any FSPF link cost, ensure that the FSPF path distance (that is, the sum of the link costs on the path) of any IVR path is less than 30,000.



### Note

IVR-enabled VSANs can be configured when the interop mode is enabled (any interop mode) or disabled (no interop mode).

## Domain ID Guidelines

Domain IDs must be unique across inter-connected VSANs when not using IVR NAT. To ensure unique domain IDs across inter-connected VSANs, consider these guidelines:

- Minimize the number of switches that require a domain ID assignment. This ensures minimum traffic disruption.
- Minimize the coordination between interconnected VSANs when configuring the SAN for the first time as well as when you add each new switch.

You can configure domain IDs using one of two options:

- Configure the allowed-domains list so that the domains in different VSANs are non-overlapping on all participating switches and VSANs.
- Configure static, non-overlapping domains for each participating switch and VSAN.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Transit VSAN Guidelines

Before configuring transit VSANS, consider the following guidelines:

- Besides defining the IVR zone membership, you can choose to specify a set of transit VSANs to provide connectivity between two edge VSANs:
  - If two edge VSANs in an IVR zone overlap, then a transit VSAN is not required (though, not prohibited) to provide connectivity.
  - If two edge VSANs in an IVR zone do not overlap, you may need one or more transit VSANs to provide connectivity. Two edge VSANs in an IVR zone will not overlap if IVR is not enabled on a switch that is a member of both the source and destination edge VSANs.
- Traffic between the edge VSANs only traverses through the shortest IVR path.
- Transit VSAN information is common to all IVR zone sets. Sometimes, a transit VSAN can also double-up as an edge VSAN in another IVR zone.

## Border Switch Guidelines

Before configuring border switches, consider the following guidelines:

- Border switches require Cisco MDS SAN-OS Release 1.3(1) or later.
- A border switch must be a member of two or more VSANs.
- A border switch that facilitates IVR communications must be IVR enabled.
- IVR can (optionally) be enabled on additional border switches to provide redundant paths between active IVR zone members.
- The VSAN topology configuration must be updated before a border switch is added or removed.

# Using the IVR Zone Wizard

The IVR Zone Wizard simplifies the steps required to configure IVR zones in a fabric. The IVR Zone Wizard checks the following conditions and prompts you for any issues:


- Checks if all switches in the fabric are Cisco MDS SAN-OS Release 2.1(1a) or later and if so, asks if you want to migrate to using IVR NAT with auto-topology.
- Checks if any switches in the fabric are earlier than Cisco MDS SAN-OS Release 2.1(1a) and if so, asks you to upgrade the necessary switches or to disable IVR NAT or auto-topology if they are enabled.

To use the IVR Zone Wizard to configure IVR and IVR zones, follow these steps:

- 
- Step 1** From Fabric Manager, click the **IVR Zone Wizard** icon in the Zone toolbar. [Figure 16-2](#) shows the IVR Zone Wizard icon.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Figure 16-2**      **IVR Zone Wizard Icon**



You see the IVR Zone Wizard.

**Step 2**      Select the VSANs that will participate in IVR in the fabric.

**Step 3**      Select the end devices that you want to communicate over IVR.



**Note**      If you are not using IVR NAT, Fabric Manager may display an error message if all the switches participating in IVR do not have unique domain IDs. You must reconfigure those switches before configuring IVR.

**Step 4**      If you enable IVR NAT, verify switches that Fabric Manager will enable with IVR NAT, CFS for IVR, and IVR topology in auto mode.



**Note**      You may have to configure a unique AFID for switches in the fabric that have non-unique VSAN IDs.

**Step 5**      If you did not enable IVR NAT, verify the transit VSAN or configure the transit VSAN if Fabric Manager cannot find an appropriate transit VSAN.

**Step 6**      Set the IVR zone and IVR zone set.

**Step 7**      Verify all steps that Fabric Manager will take to configure IVR in the fabric.

**Step 8**      Click **Finish** if you want to enable IVR NAT and IVR topology and create the associated IVR zones and IVR zone set, or click **Cancel** to exit the IVR Wizard without saving any changes.



**Note**      IVR NAT and auto-topology can be configured independently if you configure these features outside the IVR Zone Wizard. See the [“Modifying IVR” section on page 16-8](#).

## Modifying IVR

You can modify IVR using the IVR tables in the Information pane in Fabric Manager. Use these tables only if you are familiar with all IVR concepts. We recommend you configure IVR using the IVR Wizard.



**Note**      Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is clicked, the other tabs in the Information pane are activated.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Modifying IVR NAT and IVR Auto Topology

To modify IVR in NAT mode and IVR topology in auto mode from Fabric Manager, follow these steps:

- 
- |                |  |
|----------------|--|
| <b>Step 1</b>  | Select <b>All VSANs &gt; IVR</b> from the Logical Domains pane. You see the IVR configuration in the Information pane.   |
| <b>Step 2</b>  | Select the <b>CFS</b> tab if CFS is enabled for this feature in the fabric.  |
| <b>Step 3</b>  | Select <b>enable</b> from the Enable/Admin column for the primary switch.  |
| <b>Step 4</b>  | Select the <b>Apply Changes</b> button from the Information pane to distribute this change to all switches in the fabric, or select the <b>Undo Changes</b> button to cancel any changes you made. |
| <b>Step 5</b>  | Select the <b>Actions</b> tab.   |
| <b>Step 6</b>  | Check the <b>Enable IVR Nat</b> check box to enable IVR in NAT mode.   |
| <b>Step 7</b>  | Check the <b>Automatically Discover Topology</b> check box to enable IVR topology in auto mode.  |
| <b>Step 8</b>  | Select the <b>Apply Changes</b> button from the Information pane to enable IVR on the switches, or select the <b>Undo Changes</b> button to cancel any changes you made.                           |
| <b>Step 9</b>  | Click <b>CFS &gt; Config Changes &gt; Action</b> and choose <b>commit</b> .  |
| <b>Step 10</b> | Select the <b>Apply Changes</b> button from the Information pane to distribute IVR on the switches.  |
- 

## Configuring Service Group

A service group limits the scope of IVR CFS traffic across the fabric. The service group includes all IVR-enabled switches and associated VSANs.

To configure service group using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Select <b>All VSANs &gt; IVR</b> from the Logical Domains pane. You see the IVR configuration in the Information pane.     |
| <b>Step 2</b> | Select the <b>Service Group</b> tab to display the existing service groups.  |
| <b>Step 3</b> | Click the <b>Create Row</b> icon to create a new service group. You see the service group dialog box.                      |
| <b>Step 4</b> | Check the switch check box for each switch involved in IVR.  |
| <b>Step 5</b> | Set the <b>Name</b> of the service group and set the <b>Fabric ID</b> for this entry.                                      |
| <b>Step 6</b> | Enter a comma-separated list of VSAN IDs in the <b>VSAN List</b> text box.   |
| <b>Step 7</b> | Click <b>Create</b> to create this entry or click <b>Cancel</b> to discard all changes.                                    |
| <b>Step 8</b> | Repeat <a href="#">Step 1</a> through <a href="#">Step 7</a> for all switches and AFIDs associated with your IVR topology. |
- 

## Configuring AFIDs

You configure AFIDs individually for VSANs, or you set the default AFIDs for all VSANs on a switch. If you configure an individual AFID for a subset of the VSANs on a switch that has a default AFID, that subset uses the configured AFID while all other VSANs on that switch use the default AFID.

To configure default AFIDs using Fabric Manager, follow these steps:

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
  - Step 2** Select the **Default Fabric ID** tab to display the existing default AFIDs.
  - Step 3** Click the **Create Row** icon to create a default AFID. You see the default AFID dialog box.
  - Step 4** Check the switch check box for each switch involved in IVR that you want to use this default AFID.
  - Step 5** Set the **SwitchWWN** and set the default **Fabric ID** for this entry.
  - Step 6** Click **Create** to create this entry or click **Cancel** to discard all changes.
  - Step 7** Repeat [Step 1](#) through [Step 6](#) for all switches and default AFIDs you want to configure in your IVR topology.
- 

To configure individual AFIDs using Fabric Manager, follow these steps:

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
  - Step 2** Select the **Fabric ID** tab to display the existing AFIDs.
  - Step 3** Click the **Create Row** icon to create an AFID. You see the AFID dialog box.
  - Step 4** Check the switch check box for each switch involved in IVR that you want to use this default AFID.
  - Step 5** Set the **SwitchWWN** and set the **Fabric ID** for this entry.
  - Step 6** Enter a comma-separated list of VSAN IDs in the **VSAN List** text box.
  - Step 7** Click **Create** to create this entry or click **Cancel** to discard all changes.
  - Step 8** Repeat [Step 1](#) through [Step 6](#) for all switches and AFIDs you want to configure in your IVR topology.
- 

## Enabling IVR Without NAT

To enable IVR without IVR in NAT mode from Fabric Manager, follow these steps:

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
  - Step 2** Select the **CFS** tab if CFS is enabled for this feature in the fabric.
  - Step 3** Select **enable** from the Enable/Admin column for the primary switch.
  - Step 4** Select the **Apply Changes** button from the Information pane to distribute this change to all switches in the fabric, or select the **Undo Changes** button to cancel any changes you made.
  - Step 5** If CFS is not enabled, select the **Control** tab if it is not already displayed to enable IVR individually for each switch.
  - Step 6** Set the command drop-down menu to enable for each switch you want to enable IVR on.
  - Step 7** Select the **Apply Changes** button from the Information pane to enable IVR on the switches, or select the **Undo Changes** button to cancel any changes you made.
  - Step 8** Click **CFS > Config Changes > Action** and choose **commit**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 9** Select the **Apply Changes** button from the information pane to distribute IVR on the switches.
- 

## Manually Creating the IVR Topology

You must create the IVR topology in every IVR-enabled switch in the fabric if you have not configured IVR topology in auto mode. You can have up to 128 VSANs in an IVR topology. Specify the IVR topology using the following information:

- The switch WWNs of the IVR-enabled switches.
- A minimum of two VSANs to which the IVR-enabled switch belongs.
- The AFID, which distinguishes two VSANs that are logically and physically separate, but have the same VSAN number (segmented VSANs). Cisco MDS SAN-OS Release 2.1(1a) supports up to 64 AFIDs.



**Note** Two VSANs with the same VSAN number but different AFIDs are counted as two VSANs out of the total 128 VSANs allowed in the fabric.

---



**Note** The use of a single AFID does not allow for segmented VSANs in an inter-VSAN topology.

---

To create the IVR topology from Fabric Manager, follow these steps:

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
- Step 2** Select the **Local Topology** tab to display the existing IVR topology.
- Step 3** Select the **Create Row** button from the Information pane to create rows in the IVR topology. You see the local topology create dialog box.
- Step 4** Select the switch, switch WWN, and a comma-separated list of VSAN IDs for this entry.
- Step 5** Select the **Create** button to create this new row, or select **Cancel** to cancel all changes.
- Step 6** Select the **Apply Changes** button from the Information pane to create the IVR topology, or select the **Undo Changes** button to cancel any changes you made.
- 



**Note** Repeat this configuration in all IVR-enabled switches or distribute using CFS.

---



**Tip** Transit VSANs are deduced based on your configuration. The IVR feature does not have an explicit transit-VSAN configuration.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Activating an IVR Topology

After creating the IVR topology, you must activate it.

To activate the IVR topology from Fabric Manager, follow these steps:

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
  - Step 2** Select the **Action** tab to display the existing IVR topology.
  - Step 3** Select the **Activate Local** check box.
  - Step 4** Select the **Apply Changes** button from the Information pane to activate the IVR topology, or select the **Undo Changes** button to cancel any changes you made.
- 



### Caution

Active IVR topologies cannot be deactivated.

---

## Clearing the IVR Topology

You can only clear manually created IVR VSAN topology entries from the config database.

To clear the IVR topology from Fabric Manager, follow these steps:

- 
- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
  - Step 2** Select the **Control** tab if it is not already displayed.
  - Step 3** Highlight the rows you want to delete from the IVR topology.
  - Step 4** Select the **Delete Row** button from the Information pane to delete these rows from the IVR topology .
  - Step 5** Select the **Apply Changes** button from the Information pane to delete the IVR topology, or select the **Undo Changes** button to cancel any changes you made.
- 

## Adding IVR Virtual Domains

In a remote VSAN, the IVR application does not automatically add the virtual domain to the assigned domains list. Some switches (for example, the Cisco SN5428) do not query the remote name server until the remote domain appears in the assigned domains list in the fabric. In such cases, add the IVR virtual domains in a specific VSAN(s) to the assigned domains list in that VSAN. When adding IVR domains, all IVR virtual domains that are currently present in the fabric (and any virtual domain that is created in the future) will appear in the assigned domain list for that VSAN.



### Tip

Be sure to add IVR virtual domains if the following conditions apply:

- When an IVR zone set is not active.
- If Cisco SN5428 or Qlogic SANBox switches exist in the VSAN.

---



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**



**Tip**



**Tip**

As of Cisco MDS SAN-OS Release 1.3(4a), only add IVR domains in the edge VSANs and not in transit VSANs.

When you enable the IVR virtual domains, links may fail to come up due to overlapping virtual domain identifiers. If so, temporarily withdraw the overlapping virtual domain from that VSAN.



**Note**

Withdrawing an overlapping virtual domain from an IVR VSAN disrupts IVR traffic to and from that domain.

To add IVR virtual domains using Fabric Manager, follow these steps:

- Step 1** Select **All VSANs > IVR** from the Logical Domains pane. You see the IVR configuration in the Information pane.
- Step 2** Select the **Action** tab to display the existing IVR topology.
- Step 3** Enter a comma-separated list of VSAN IDs in the **Create Virtual Domain for VSANs** column.
- Step 4** Select the **Apply Changes** button from the Information pane to activate the IVR topology, or select the **Undo Changes** button to cancel any changes you made.

## IVR Zones and IVR Zone Sets

As part of the IVR configuration, you need to configure one or more IVR zone sets to enable cross-VSAN communication. To achieve this result, you must specify each IVR zone as a set of (pWWN, VSAN) entries. Like zones, several IVR zone sets can be configured to belong to an IVR zone. You can define several IVR zone sets and activate only one of the defined IVR zone sets.



**Note**

The same IVR zone set must be activated on *all* of the IVR-enabled switches.

## IVR Zones Versus Zones

Table 16-1 identifies the key differences between IVR zones and zones.

**Table 16-1 Key Differences Between IVR Zones and Zones**

IVR Zones	Zones
IVR zone membership is specified using the VSAN and pWWN combination.	Zone membership is specified using pWWN, fabric WWN, sWWN, or the AFID.
Default zone policy is always deny (not configurable).	Default zone policy is deny (configurable).

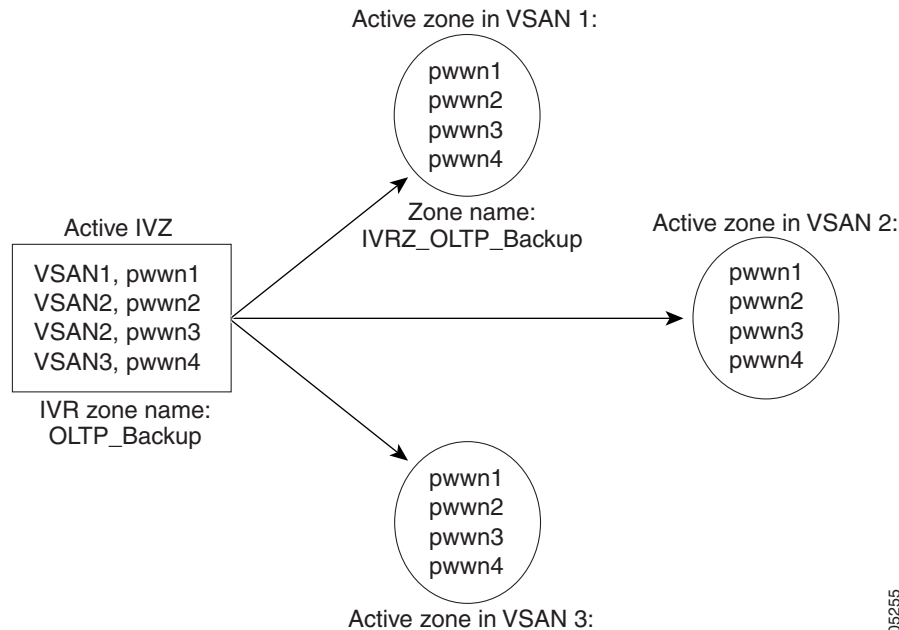
**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Automatic IVR Zone Creation

Figure 16-3 depicts an IVR zone consisting of four members. To allow pwwn1 to communicate with pwwn2, they must be in the same zone in VSAN 1, as well as in VSAN 2. If they are not in the same zone, then the hard-zoning ACL entries will prohibit pwwn1 from communicating with pwwn2.

A zone corresponding to each active IVR zone is automatically created in each edge VSAN specified in the active IVR zone. All pWWNs in the IVR zone are members of these zones in each VSAN.

**Figure 16-3**      *Creating Zones upon IVR Zone Activation*



The zones are created automatically by the IVR process when an IVR zone set is activated. They are not stored in a full zone set database and are lost when the switch reboots or when a new zone set is activated. The IVR feature monitors these events and adds the zones corresponding to the active IVR zone set configuration when a new zone set is activated. Like zone sets, IVR zone sets are also activated nondisruptively.



### Note

If pwwn1 and pwwn2 are in an IVR zone in the current as well as the new IVR zone set, then activation of the new IVR zone set does not cause any traffic disruption between them.

IVR zone and IVR zone set names are restricted to 64 alphanumeric characters.

## Configuring IVR Zones and Zone Sets

To create IVR zones or zone sets using Fabric Manager, follow these steps:

- Step 1**      Select the VSAN that you want to configure from the Logical Domains tree.
- Step 2**      Choose **Zone > IVR > Edit Local Full Zone Database** from the Zone menu.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You see the Edit IVR Local Full Zone Database dialog box for the VSAN you selected.

**Step 3** Right-click the zone set or zone for that VSAN and select **Insert** to add a zone set or zone.

If you are adding a zone set, you can activate it by right-clicking the newly created zone set and selecting **Activate**. This configuration is distributed to the other switches in the network fabric.



**Note**

When you confirm the activate operation, the current running configuration is saved to the startup configuration. This permanently saves any changes made to the running configuration (not just zoning changes).



**Note**

Sometimes zones beginning with prefix IVRZ and a zone set with name nozoneset appear in logical view. The zones with prefix IVRZ are IVR zones that get appended to regular active zones. The prefix IVRZ is appended to active IVR zones by the system. Similarly the zone set with name nozoneset is an IVR active zone set created by the system if no active zone set is available for that VSAN and if the `ivrZonesetActivateForce` flag is enabled on the switch. In the `server.properties` file, you can set the property `zone.ignoreIVRZones` to true or false to either hide or view IVR zones as part of regular active zones. See the [“Fabric Manager Server Properties File” section on page 2-8](#) for more information on the `server.properties` file.



**Note**

Do not create a zone with prefix the IVRZ or a zone set with name nozoneset. These names are used by the system for identifying IVR zones.

**Step 4** Optionally, check the **Permit QoS Traffic with Priority** check box and set the QoS priority for this zone.

**Step 5** Click **OK** to create this zone or zone set or click **Close** to discard all changes.

## Creating Additional IVR Zones and Zone Sets

To create additional zones and zone sets using Fabric Manager, follow these steps:

**Step 1** Click **Zone > IVR > Edit Local Full Zone Database**. You see the Edit Local Full Zone Database dialog box.

**Step 2** Right-click the **Zones** folder and choose **Insert** from the pop-up menu.

**Step 3** Enter the zone name in the dialog box that appears and click **OK** to add the zone.

The zone is automatically added to the zone database.

**Step 4** Right-click the **ZoneSets** folder in the Edit Local Full Zone Database dialog box, and choose **Insert**.

**Step 5** Enter the zone set name in the dialog box that appears and click **OK** to add the zone set.

The zone set is automatically added to the zone database.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Activating IVR Zone Sets

Once the zone sets have been created and populated, you must activate the zone set.

To activate an IVR zone set, follow these steps:

- 
- Step 1** Click **Zone > IVR > Edit Local Full Zone Database**. You see the Edit Local Full Zone Database dialog box.
  - Step 2** Right-click the **Zoneset** folder and choose the zone set that you want to activate from the pop-up menu.
  - Step 3** Click **Activate**.



**Note** The active zone set in Edit Zone is always shown in bold, even after successful activation. This is because a member of this VSAN must be participating in IVR zoning. Because the IVR zones get added to active zones, the active zone set configuration is different from the local zone set configuration with same name.

---

## Deactivating IVR Zone Sets

To deactivate a zone set, follow these steps:

- 
- Step 1** Click **Zone > IVR > Edit Local Full Zone Database**. You see the Edit Local Full Zone Database dialog box.
  - Step 2** Right-click the **Zoneset** folder and choose the zone set that you want to deactivate from the pop-up menu.
  - Step 3** Click **Deactivate**.
- 

## Recovering an IVR Full Zone Database

You can recover an IVR zone database by copying the IVR full zone database.

To recover an IVR zone database, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Zone > IVR > Copy Full Zone Database** from the Zone menu. You see the Copy Full Zone Database dialog box.
  - Step 2** Select the **Active** or the **Full** radio button, depending on which type of IVR database you want to copy.
  - Step 3** Select the source switch from which to copy the information from the drop-down list.
  - Step 4** Select the destination switch from the drop-down list.
  - Step 5** Click **Copy** to copy the database, or click **Close** to close the dialog box without copying.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Recovering an IVR Full Topology

You can recover a topology by copying from the active zone database or the full zone database.

To recover a zone topology, follow these steps.

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | From Fabric Manager, choose <b>Zone &gt; IVR &gt; Copy Full Topology</b> . You see the Copy Full Topology dialog box.    |
| <b>Step 2</b> | Select the <b>Active</b> or the <b>Full</b> radio button, depending on which type of IVR database you want to copy from. |
| <b>Step 3</b> | Select the source switch from which to copy the information from the drop-down list.                                     |
| <b>Step 4</b> | Select the destination switch from the drop-down list.   |
| <b>Step 5</b> | Click <b>Copy</b> to copy the topology, or click <b>Close</b> to close the dialog box without copying.                   |
- 

## Adding Members to IVR Zones

You can add members to existing IVR zones using the Edit Local Full Zone Database dialog box. LUN-zoning can optionally be used between members of active IVR zones.

To add members to an existing IVR zone and optionally configure LUN zoning using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Click <b>Zone &gt; IVR &gt; Edit Local Full Zone Database</b> . You see the Edit Local Full Zone Database dialog box.                   |
| <b>Step 2</b> | Expand the <b>Zones</b> folder and choose the zone you want to add a member to.   |
| <b>Step 3</b> | Click the <b>Insert</b> icon to add a new member in this zone. You see the zone membership dialog box.                                  |
| <b>Step 4</b> | Set the <b>WWN</b> for the end device you want to add as a member of this IVR zone.   |
| <b>Step 5</b> | Set the <b>Port VSAN Id</b> and <b>Fabric ID</b> for this end device.   |
| <b>Step 6</b> | Optionally, check the <b>LUNs</b> check box and set the LUNs you want this IVR zone to access on this end device.                       |
| <b>Step 7</b> | Click <b>Add</b> to add the member to the IVR zone with the optional LUN zoning attribute or click <b>Close</b> to discard all changes. |
- 

## IVR Interoperability

When using the IVR feature, all border switches in a given fabric must be Cisco MDS switches. However, other switches in the fabric may be non-MDS switches. For example, end devices that are members of the active IVR zone set may be connected to non-MDS switches. Non-MDS switches may also be present in the transit VSAN(s) or in the edge (VSANs) if one of the **interop** modes is enabled.

See the “[Switch Interoperability](#)” section on page 23-4.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## PortChannel Configuration

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

This chapter describes how to use the PortChannel wizard to configure PortChannels. This chapter contains the following sections:

- [PortChannel Functionality, page 17-1](#)
- [Using the PortChannel Wizard, page 17-2](#)
- [Modifying PortChannels, page 17-5](#)

## PortChannel Functionality

A PortChannel has the following functionality:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).
- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

PortChanneling and trunking are used separately across an ISL:

- PortChanneling, which enables several links to be combined into one aggregated link, can be done between E ports and TE ports.
- Trunking, which permits carrying traffic on multiple VSANs between switches, can be done only between TE ports.

The Cisco MDS 9000 Family of switches supports 128 PortChannels with 16 interfaces per PortChannel. A PortChannel number refers to the unique (to each switch) identifier associated with each channel group. This number ranges from 1 to 128.


**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Using the PortChannel Wizard

To create a PortChannel with the PortChannel Wizard in Fabric Manager, follow these steps:

- Step 1** Click the **PortChannel Wizard** icon in the toolbar (see [Figure 17-1](#)).  
You see the first PortChannel Wizard screen.

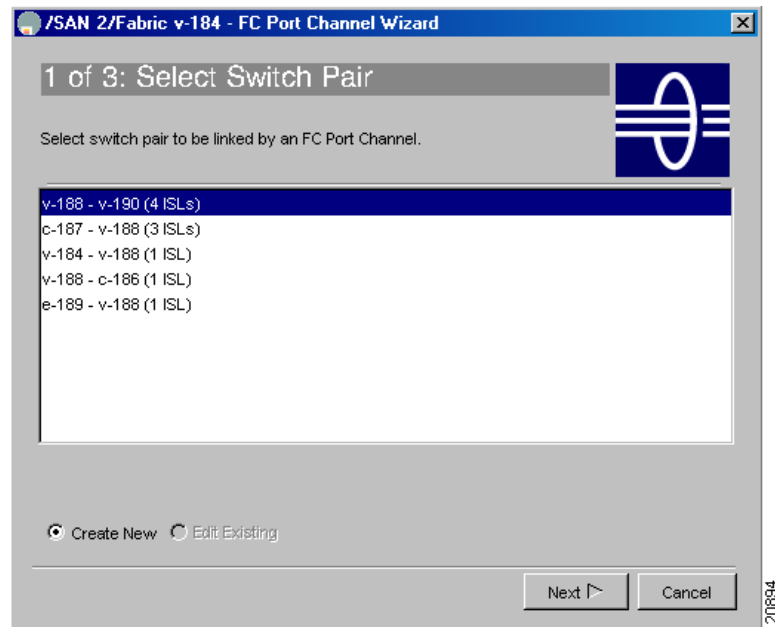
**Figure 17-1** PortChannel Wizard Icon



PortChannel Wizard icon

- Step 2** Select a switch pair. [Figure 17-2](#) shows a list of the switch pairs.

**Figure 17-2** Select Switch Pairs




- Step 3** Click **Next**.  
**Step 4** Select the ISLs.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Figure 17-3 shows a list of the ISLs.

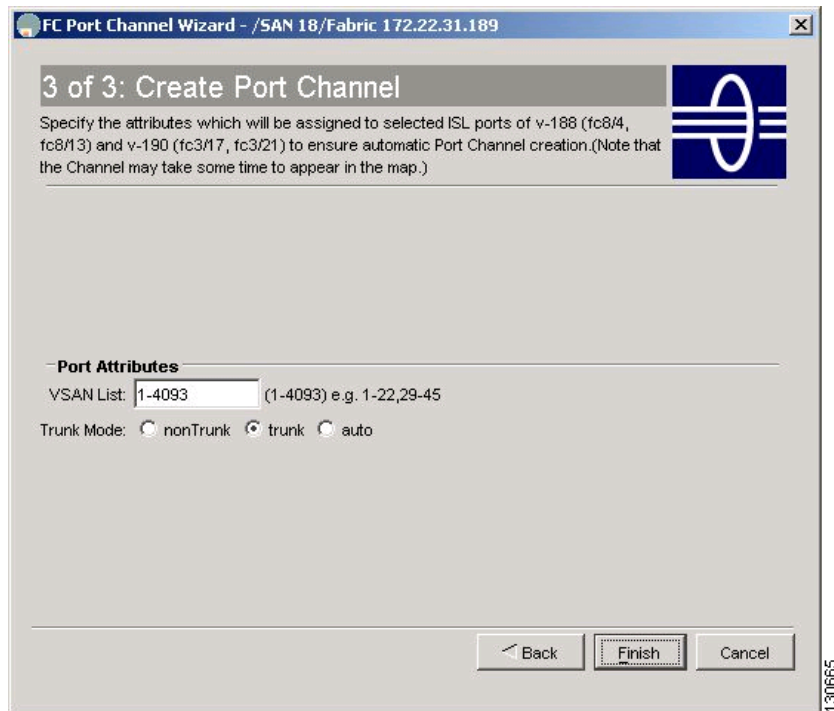
**Figure 17-3**      **Select ISLs**



- Step 5**      Optionally, check the **Dynamically form Port Channel Group from selected ISLs** check box if you want to dynamically create the PortChannel and make the ISL properties identical for Admin, Trunk, Speed, and VSAN attributes.
- Step 6**      Click **Next**.
- Step 7**      If you chose to dynamically form a portChannel from selected ISLs, you see the the final PortChannel Wizard screen as shown in Figure 17-4. Set the **VSAN List** and **TrunkMode** and proceed to Step 11.

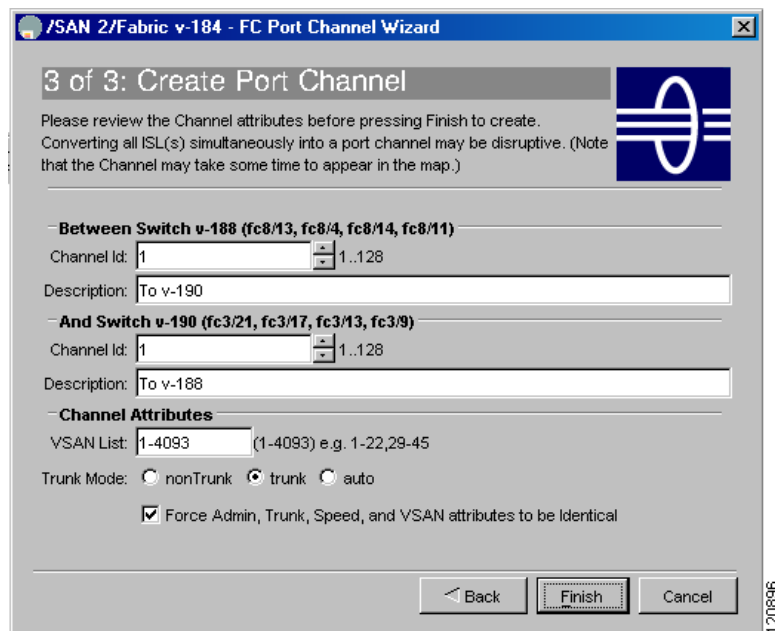
**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 17-4** Dynamically Form a PortChannel



**Step 8** If you did not choose to dynamically form a PortChannel, you see the third PortChannel Wizard screen as shown in [Figure 17-5](#).

**Figure 17-5** Create a PortChannel



**Step 9** Change the channel ID or description for each switch, if necessary.

**Step 10** Review the attributes at the bottom of the screen, and set them if applicable.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The following attributes are shown in [Figure 17-5](#):

- **VSAN List**—The list of VSANs to which the ISLs belong.
- **Trunk Mode**—You can enable trunking on the links in the PortChannel. Select **trunking** if your link is between TE ports. Select **nontrunking** if your link is between E ports (for example, if your link is between an MDS switch and another vendor's switch). Select **auto** if you are not sure.
- **Force Admin, Trunk, Speed, and VSAN attributes to be identical**—This ensures that the same parameter settings are used in all physical ports in the channel. If these settings are not identical, the ports cannot become part of the PortChannel.

**Step 11** Click **OK**.

The PortChannel is created. Note that it may take a few minutes before the new PortChannel is visible in the Fabric pane.

---

## Modifying PortChannels

To modify an existing PortChannel configuration using Fabric Manager, follow these steps:

- 
- Step 1** Choose **ISLs > Port Channels** from the Physical Attributes Pane. You see the PortChannels configured in the Information pane.
- Step 2** Choose the **Channels** tab to add a PortChannel or modify an existing PortChannel.
- Step 3** Choose the **Protocols** tab to change the mode for an existing PortChannel.
- Step 4** Click the **Apply Changes** icon to save any modifications or click **Undo Changes** to discard any changes.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Interface Configuration

---

A switch's main function is to relay frames from one data link to another. To do that, the characteristics of the interfaces through which the frames are received and sent must be defined. This chapter describes the basic interface configuration to get your switch up and running.

This chapter includes the following sections:

- [Fibre Channel Interfaces, page 18-1](#)
- [Configuring Fibre Channel Interfaces, page 18-6](#)
- [Enabling or Disabling Interfaces, page 18-7](#)
- [Managing Interface Attributes for Ports, page 18-7](#)
- [Configuring the Management Interface, page 18-9](#)
- [IPFC Interface Configuration, page 18-9](#)

## Fibre Channel Interfaces

This section describes Fibre Channel interface characteristics, including (but not limited to) modes, states, and speeds. It includes the following sections:

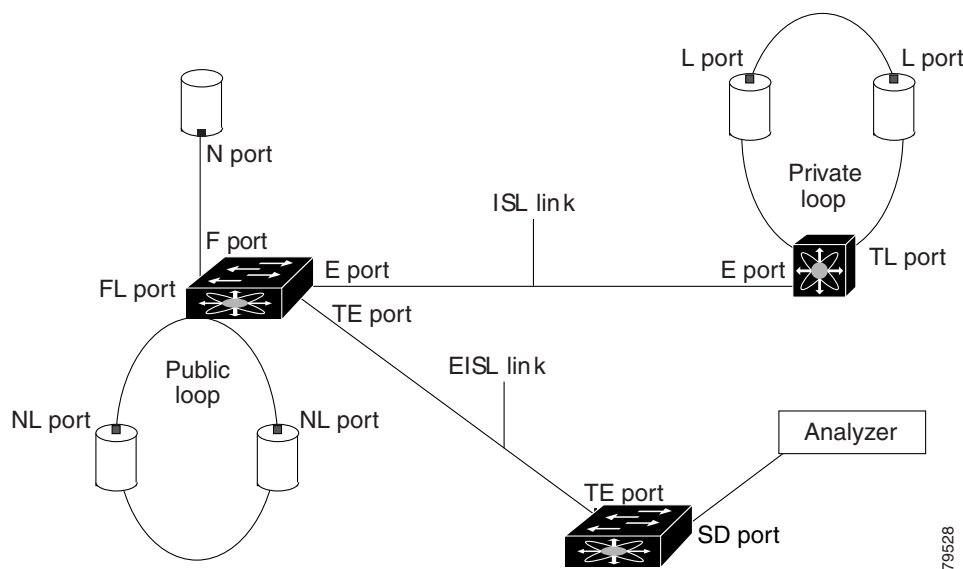
- [About Interface Modes, page 18-1](#)
- [About Interface States, page 18-5](#)

## About Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several modes: E port, F port, FL port, TL port, TE port, SD port, ST port, and B port (see [Figure 18-1](#)). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 18-1 Cisco MDS 9000 Family Switch Interface Modes**



**Note**

Interfaces are created in VSAN 1 by default. See [Chapter 13, “VSAN Configuration.”](#)

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.
- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

A brief description of each interface mode follows.

## E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see [Chapter 17, “PortChannel Configuration”](#)).

## F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

## TL Port

In translatable loop port (TL port) mode, an interface functions as a translatable loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop.



### Tip

---

We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

---

## TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an Extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family. TE ports support class 2, class 3, and class F service.

## SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes through a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they merely transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## ST Port

In the SPAN Tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic.

## Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

## B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

## Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, or TE port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode.

TL ports and SD ports are not determined during initialization and are administratively configured.

## Configuring Trunking Mode

To configure trunking mode on an interface, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Interfaces &gt; FC Physical</b> in Fabric Manager. You see the interface configuration in the Information pane.   |
| <b>Step 2</b> | Right-click the interface you want to configure in the Device view on Device Manager and choose <b>Configure....</b> You see the interface configuration dialog box.  |
| <b>Step 3</b> | Choose the <b>Trunk Config</b> tab to modify the trunking mode for the selected interface.  |
| <b>Step 4</b> | Click <b>Apply Changes</b> on Fabric Manager or <b>Apply</b> on Device Manager to save these changes or click <b>Undo Changes</b> on Fabric Manager or <b>Close</b> on Device Manager to discard any unsaved changes. |
- 

**Tip**

Configure one side of the trunking link as trunk and the other side as auto for best results.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## About Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

### Administrative States

The administrative state refers to the administrative configuration of the interface as described in [Table 18-1](#).

**Table 18-1** Administrative States

Administrative State	Description
Up	Enables an interface.
Down	Disables an interface. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

### Operational States

The operational state indicates the current operational state of the interface as described in [Table 18-2](#).

**Table 18-2** Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed.
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE mode.

### Reason Codes

Reason codes are dependent on the operational state of the interface as described in [Table B-1](#).

## 32-Port Configuration Guidelines

The 32-port guidelines applies to the following hardware:

- The 32-port 2 Gbps or 1 Gbps switching module
- The Cisco MDS 9140 Switch

When configuring these host-optimized ports, the following port mode guidelines apply:

You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8 and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8 and so on) are not usable and remain shutdown.

- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- The auto mode is the default port mode. The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules and the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).

**Note**

In the Cisco MDS 9100 Series, the left most groups of ports outlined in white (4 ports in the 9120 switch and 8 ports in the 9140 switch) are full line rate like the 16-port switching module. The other ports (16 ports in the 9120 switch and 32 ports in the 9140 switch) are host-optimized like the 32-port switching module. Each group of 4 host-optimized ports have the same rules as for the 32-port switching module. Use the Device Manager Show Oversubscription dialog box to monitor these ports.

## Configuring Fibre Channel Interfaces

To configure Fibre Channel port interfaces, follow these steps:

- Step 1** Choose **Switches > Interfaces > FC Physical** on Fabric Manager or right-click the interface and choose **Configure...** on Device Manager. You see the interface configuration in the Information pane on Fabric Manager or the Interface dialog box in Device Manager.
- Step 2** Choose the **General** tab and set the interface mode, port VSAN membership, and administrative state.
- Step 3** Optionally, set other configuration parameters using the other tabs.
- Step 4** Click **Apply**.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Configuring Gigabit Ethernet Interfaces

Each port or interface on the IPS module is displayed in the Ethernet Port dialog box.

To configure Ethernet port interfaces, follow these steps :

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Switches &gt; Interfaces &gt; Gigabit Ethernet</b> on Fabric Manager or right-click the interface and choose <b>Configure...</b> on Device Manager. You see the interface configuration in the Information pane on Fabric Manager or the Interface dialog box in Device Manager. |
| <b>Step 2</b> | Choose the <b>General</b> tab and set the description, status, and IP address.   |
| <b>Step 3</b> | Optionally, set other configuration parameters using the other tabs.   |
| <b>Step 4</b> | Click <b>Apply</b> .   |
- 

## Enabling or Disabling Interfaces

To enable an interface using Device manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Right-click an interface and choose <b>Enable</b> or <b>Disable</b> from the pop-up menu.                         |
| <b>Step 2</b> | To enable or disable multiple interfaces, Ctrl-click each port or drag the mouse around a group of interfaces.    |
| <b>Step 3</b> | Right-click any of the selected interfaces and click either <b>Enable</b> or <b>Disable</b> from the pop-up menu. |
- 

To enable an interface using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; Interfaces &gt; Gigabit Ethernet</b> or <b>Switches &gt; Interfaces &gt; FC Physical</b> . You see the interface configuration in the Information pane. |
| <b>Step 2</b> | Choose the <b>General</b> tab and set <b>Status &gt; Admin</b> to up (for enable) or down (for disable).  |
| <b>Step 3</b> | Optionally, set other configuration parameters using the other tabs.  |
| <b>Step 4</b> | Click the <b>Apply Changes</b> icon.  |
- 

## Managing Interface Attributes for Ports

To manage interface attributes from the Fabric Manager, choose **Switches > Interfaces** from the Physical tree in the Navigation pane and then choose one of the following types to be configured:

- FC Physical
- FC Logical

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Ethernet
- SVC
- PortChannels

To manage interface attributes from the Device Manager, right-click a port on a module, and then click **Configure** from the pop-up menu or choose an interface type from the Interface menu.

The Fabric Manager Information pane displays interface attributes for multiple switches. The dialog box from Device Manager displays interface attributes for a single switch.

## Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB\_credits) are a flow control mechanism to ensure that FC switches do not run out of buffers, because switches must not drop frames. BB\_credits are negotiated on a per-hop basis.

The receive BB\_credit (rxbbcredit) value may be configured for each FC interface. In most cases, you do not need to modify the default configuration.



### Note

The receive BB\_credit values depend on the module type and the port mode:

- 16-port switching modules and full rate ports: The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
- 32-port switching modules and host-optimized ports: The default value is 12 for the Fx, E, and TE modes. These values cannot be changed.

## Performance Buffers

Regardless of the configured Rx BB\_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB\_credit value.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

## Configuring Buffer-to-Buffer Credits and Performance Buffers

To configure buffer to buffer credits or performance buffers, follow these steps:

- Step 1** Choose **Switches > Interfaces > FC Physical** on Fabric Manager or right-click the interface and choose **Configure...** on Device Manager. You see the interface configuration in the Information pane on Fabric Manager or the Interface dialog box in Device Manager.
- Step 2** Choose the **BB Credit** tab and set the buffer-to-buffer credits or performance buffers for the selected interface.
- Step 3** Click **Apply**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Identification of SFP Types

To show the SFP types for an interface using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Switches &gt; Interfaces &gt; FC Physical</b> . You see the interface configuration in the Information pane on Fabric Manager. |
| <b>Step 2</b> | Choose the <b>Physical</b> tab to see the transmitter type for the selected interface.   |
- 

## Configuring the Management Interface

For information on configuring the management interface, refer to the *Cisco MDS 9000 Configuration Guide*.

## Configuring Persistent FC IDs

To configure persistent FC IDs using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Sanxx &gt; VSANxx &gt; Domain Manager</b> . You see the domain manager configuration in the Information pane. |
| <b>Step 2</b> | Choose the <b>Persistent FC ID</b> tab to configure persistent FC IDs.  |
| <b>Step 3</b> | Click the <b>Apply Changes</b> icon to save these FC IDs or click <b>Undo Changes</b> to discard any unsaved changes.   |
- 

## IPFC Interface Configuration

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. IPFC interfaces cannot be created for nonexistent VSANs.

Follow these guidelines when creating or deleting IPFC interfaces:

- Create a VSAN before creating the interface for that IPFC interface. If a VSAN does not exist, the interface cannot be created.
- Create the IPFC interface—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.



### Tip

---

After configuring the FCIP interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) features.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## FCIP Configuration

---

Cisco MDS 9000 Family IP storage services (IPS) modules extend the reach of Fibre Channel SANs by using open standard, IP-based technology. The switch connects separated SAN islands using Fibre Channel over IP (FCIP).



**Note**

FCIP features are specific to the IPS modules running Cisco MDS SAN-OS Release 1.1(x) or later and the Gigabit Ethernet ports on the MPS-14/2 module running Cisco MDS SAN-OS Release 2.0(x) or later.

This chapter includes the following sections:

- [About Gigabit Ethernet Interfaces, page 19-1](#)
- [FCIP Configuration, page 19-2](#)
- [FCIP Write Acceleration, page 19-4](#)
- [FCIP Compression, page 19-5](#)
- [Using the FCIP Wizard, page 19-5](#)
- [Modifying FCIP Links, page 19-8](#)
- [FCIP Tape Acceleration, page 19-11](#)
- [Configuring Advanced FCIP Interfaces, page 19-13](#)
- [FCIP High Availability, page 19-18](#)

## About Gigabit Ethernet Interfaces

Both FCIP and iSCSI rely on TCP/IP for network connectivity. On each IPS module, connectivity is provided in the form of Gigabit Ethernet interfaces that are appropriately configured. This section covers the steps required to configure IP for subsequent use by FCIP and iSCSI.

A new port mode, called IPS, is defined for Gigabit Ethernet ports on each IPS module. IP storage ports are implicitly set to IPS mode, so it can only be used to perform iSCSI and FCIP storage functions. IP storage ports do not bridge Ethernet frames or route other IP packets.



**Tip**

Gigabit Ethernet ports on any IP storage services module should not be configured in the same Ethernet broadcast domain as the management Ethernet port—they should be configured in a different broadcast domain, either by using separate standalone hubs or switches or by using separate VLANs.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring a Basic Gigabit Ethernet Interface


To configure the general characteristics of a Gigabit Ethernet interface, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Interfaces > Gigabit Ethernet** in the Physical Attributes pane. You see the Gigabit Ethernet configuration in the Information pane.
- From Device Manager, right-click the Gigabit Ethernet port that you want to configure and choose **Configure....** You see the Gigabit Ethernet configuration dialog box.
- Step 2** Click the **General** tab in Fabric Manager, or click the **GigE** tab in Device Manager to display the general configuration options for the interface.
- Step 3** Set the description and MTU value for the interface. The valid value for the MTU field can be a number in the range from 576 to 9000.
- Step 4** Set **Admin** up or down and check the **CDP** check box if you want this interface to participate in CDP.
- Step 5** Set **IpAddress/Mask** with the IP address and subnet mask for this interface.
- Step 6** From Fabric Manager, click the **Apply Changes** icon to save these changes, or click the **Undo Changes** icon to discard changes.
- From Device Manager, click **Apply** to save these changes, or click **Close** to discard changes and close the Gigabit Ethernet configuration dialog box.
- 

## FCIP Configuration

The Fibre Channel over IP Protocol (FCIP) is a tunneling protocol that connects geographically distributed Fibre Channel storage area networks (SAN islands) transparently over IP local area networks (LANs), metropolitan area networks (MANs), and wide area networks (WANs). See [Figure 19-1](#).

**Figure 19-1** Fibre Channel SANs Connected by FCIP



FCIP uses TCP as a network layer transport.

### FCIP and VE Ports

[Figure 19-2](#) describes the internal model of FCIP with respect to Fibre Channel Inter-Switch Links (ISLs) and Cisco's enhanced ISLs (EISLs).

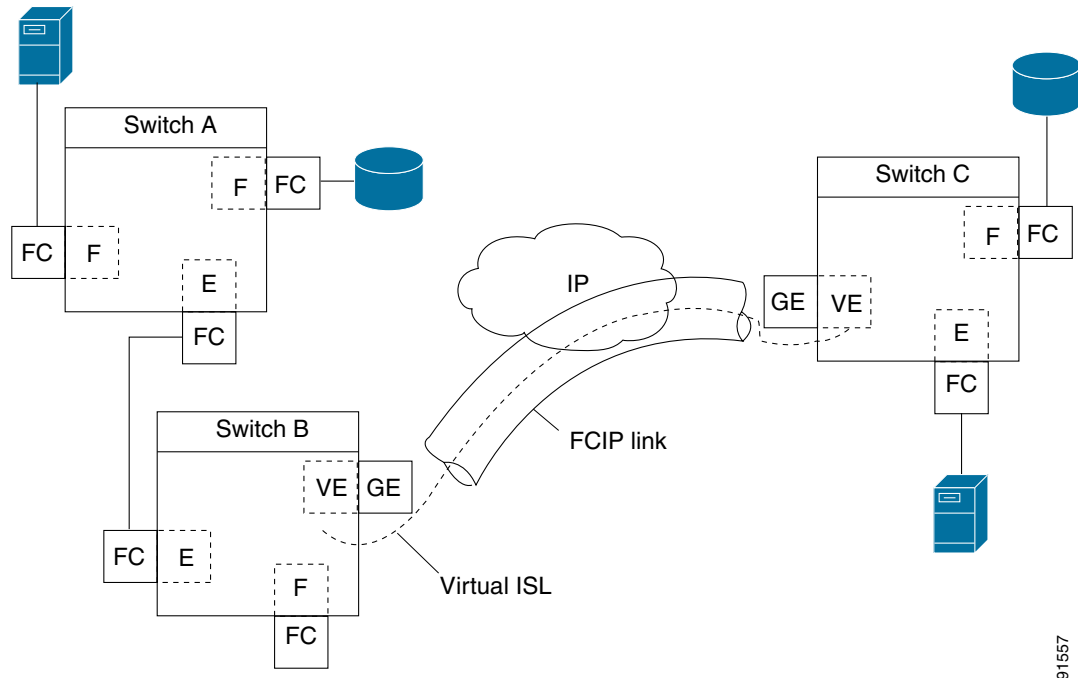
FCIP virtual E (VE) ports behave exactly like standard Fibre Channel E ports, except that the transport in this case is FCIP instead of Fibre Channel. The only requirement is that the other end of the FCIP link be another VE port.



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

A virtual ISL is established over a FCIP link and transports Fibre Channel traffic. Each associated virtual ISL looks like a Fibre Channel ISL with either an E port or a TE port at each end (see [Figure 19-2](#)).

**Figure 19-2 FCIP Links and Virtual ISLs**



See the “E Port” section on page 18-2.

## FCIP Links

FCIP links consist of one or more TCP connections between two FCIP link endpoints. Each link carries encapsulated Fibre Channel frames.

When the FCIP link comes up, the VE ports at both ends of the FCIP link create a virtual Fibre Channel (E)ISL and initiate the E port protocol to bring up the (E)ISL.

By default, the FCIP feature on any Cisco MDS 9000 Family switch creates two TCP connections for each FCIP link.

- One connection is used for data frames.
- The other connection is used only for Fibre Channel control frames, that is, switch-to-switch protocol frames (all Class F). This arrangement provides low latency for all control frames.

To enable FCIP on the IP services modules, an FCIP profile and FCIP interface (interface FCIP) must be configured.

The FCIP link is established between two peers, the VE port initialization behavior is identical to a normal E port. This behavior is independent of the link being FCIP or pure Fibre Channel and is based on the E port discovery process (ELP, ESC).

Once the FCIP link is established, the VE port behavior is identical to E port behavior for all inter-switch communication (including domain management, zones, and VSANs). At the Fibre Channel layer, VE port, and E port operations are identical.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## FCIP Write Acceleration

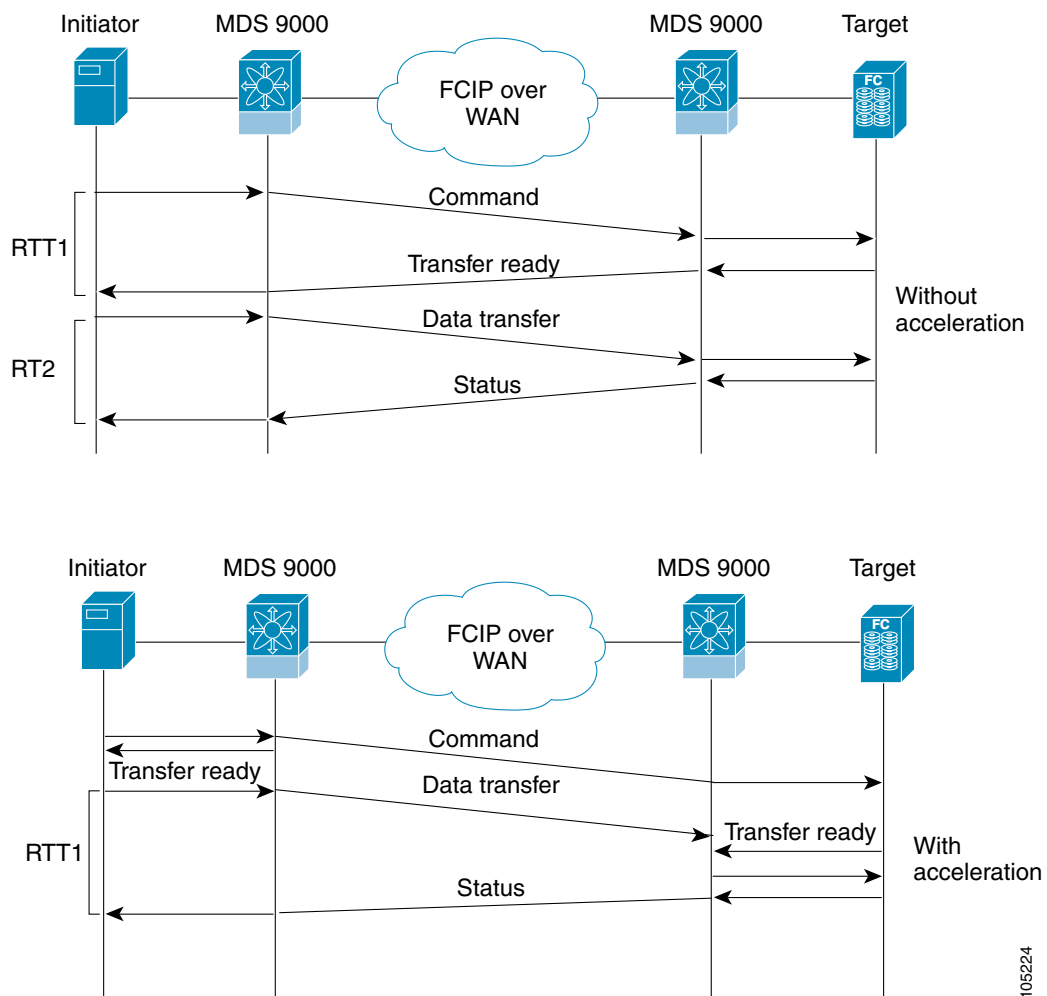
The FCIP write acceleration feature in Cisco MDS SAN-OS Release 1.3(3) enables you to significantly improve application performance when storage traffic is routed over wide area networks using FCIP. When FCIP write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for the command to transfer ready acknowledgments (see Figure 19-3).



### Note

The write acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tunnel is not initialized.

**Figure 19-3** FCIP Link Write Acceleration



In Figure 19-3, the **write** command without write acceleration requires two round-trip transfers (RTT), while the **write** command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP link back to the host before the write command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP link of the write command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP link.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Tip**

FCIP write acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.

**Caution**

When write acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if a FCIP interface is up in a FICON VSAN, write acceleration cannot be enabled on that interface.

## FCIP Compression

The FCIP compression feature introduced in Cisco MDS SAN-OS Release 1.3(x) allows IP packets to be compressed on the FCIP link if this feature is enabled on that link. By default the FCIP compression is disabled.

This feature uses the Lempel-Zif-Stac (LZS) compression algorithm to compress packets.

The high-throughput mode allows faster compression but the compression ratio may be lower. The high-comp-ratio mode allows a higher compression ratio, but the throughput may be lower.

FCIP Compression is an optional check box within the FCIP Wizard.

## Using the FCIP Wizard

**Note**

In Cisco MDS SAN-OS Release 2.0 and later, there is an additional login prompt to log into a switch that is not a part of your existing fabric.

To create and manage FCIP links with Fabric Manager, use the FCIP Wizard. First verify that the IP services module is inserted in the required Cisco MDS 9000 Family switches and that the Gigabit Ethernet interfaces on these switches are connected and the connectivity verified. The steps in creating FCIP links using the FCIP Wizard are:

- Select the endpoints.
- Choose the interfaces' IP addresses.
- Specify link attributes.
- Optionally enable FCIP write acceleration or FCIP compression.


To create FCIP links using the FCIP Wizard, follow these steps:

**Step 1**

Open the FCIP Wizard by clicking its icon in the Fabric Manager toolbar. [Figure 19-4](#) shows the FCIP Wizard icon.


**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 19-4 FCIP Wizard**



- Step 2** Choose the switches that act as endpoints for the FCIP link and click **Next**.
- Step 3** Choose the Gigabit Ethernet ports on each switch that will form the FCIP link.
- Step 4** If both Gigabit Ethernet ports are part of MPS-14/2 modules, you can check the **Enforce IPSEC Security** check box and set the **IKE Auth Key**, as shown in [Figure 19-5](#). See the “[Configuring IPsec Network Security](#)” section on [page 28-1](#) for information on IPsec and IKE.

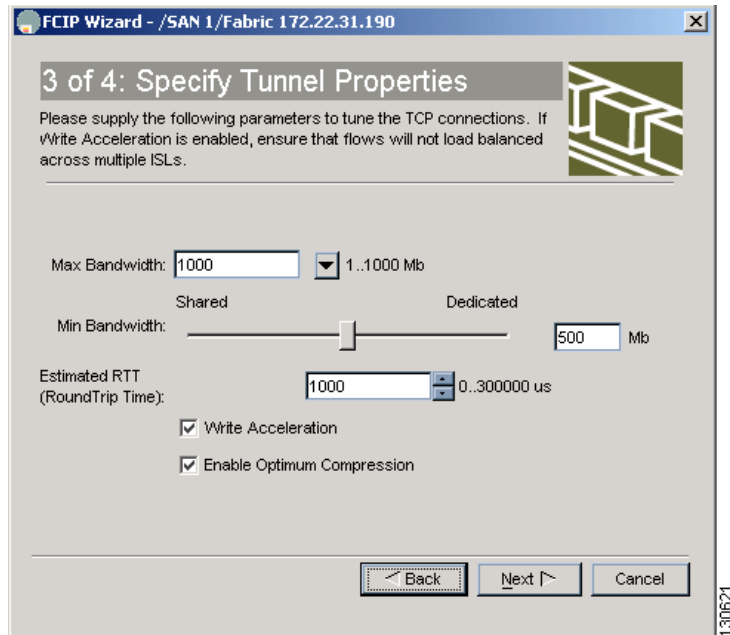
**Figure 19-5 Enabling IPsec on an FCIP link**



- Step 5** Click **Next**. You see the TCP connection characteristics.
- Step 6** Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link, as shown in [Figure 19-6](#). You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 19-6 Specifying Tunnel Properties**



- Step 7** Check the **Enable Write Acceleration** check box to enable FCIP write acceleration on this FCIP link. See the “[FCIP Write Acceleration](#)” section on page 19-4.
- Step 8** Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the “[FCIP Compression](#)” section on page 19-5.
- Step 9** Click **Next** to configure the FCIP tunnel parameters.
- Step 10** Set the **FICON Port Address** if FICON is required on this FCIP link. Click the ... button to show the first available FICON port.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

- Step 11** Set the **Port VSAN** and click the **Trunk Mode** radio button for this FCIP link, as shown in Figure 19-7. See the “Checking Trunk Status” section on page 19-10.

**Figure 19-7 Create FCIP ISL**

FCIP Wizard - /SAN 1/Fabric 172.22.31.190

**4 of 4: Create FCIP ISL**

Please supply following parameters to create a FCIP tunnel. Specify Port Vsan for nontrunk/auto and allowed Vsan list for Trunk tunnel. (Note that the FCIP link takes some time to appear in map after creation.) Security Policy with ipsec\_default\_transform\_set will be applied.

**Between Switch 172.22.31.185 (fcip2 over gigE3/1)**

IP Address/Mask:  e.g. 10.1.1.1/24

FICON Port Address:  ... 0xe0..0xf9

**And Switch 172.22.31.190 (fcip2 over gigE2/1)**

IP Address/Mask:  21.2.3.91/24 e.g. 10.1.1.1/24

FICON Port Address:  ... 0xe0..0xf9

**Attributes**

Port VSAN:  1 1..4093

Trunk Mode: ☒ nonTrunk ☐ trunk ☐ auto

- Step 12** Click **Finish** to create this FCIP link or click **Cancel** to exit the FCIP Wizard without creating an FCIP link.

## Modifying FCIP Links

Once you have created FCIP links using the FCIP wizard, you may need to modify parameters for these links. This includes modifying the FCIP profiles as well as the FCIP link parameters. Each Gigabit Ethernet interface can have three active FCIP links at one time.

## About FCIP Profiles

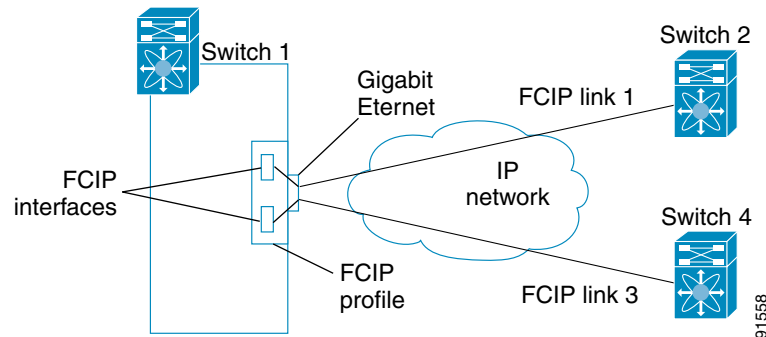
The FCIP profile contains information about local IP address and TCP parameters. The profile defines the following information:

- The local connection points (IP address and TCP port number).
- The behavior of the underlying TCP connections for all FCIP links that use this profile.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

The FCIP profile's local IP address determines the Gigabit Ethernet port where the FCIP links terminate (see [Figure 19-8](#)).

**Figure 19-8 FCIP Profile and FCIP Links**



## FCIP Interfaces

The FCIP interface is the local endpoint of the FCIP link and a VE port interface. All the FCIP and E port parameters are configured in context to the FCIP interface.

The FCIP parameters consist of the following:

- The FCIP profile determines which Gigabit Ethernet port initiates the FCIP links and defines the TCP connection behavior.
- Peer information.
- Number of TCP connections for the FCIP link.
- E port parameters—trunking mode and trunk allowed VSAN list.

## Modifying FCIP Profiles and FCIP Links

You can use Fabric Manager or Device Manager to modify FCIP links between switches or create new FCIP links. First, you must create or modify FCIP profiles, and then bind the interfaces to the profile. To bind an FCIP profile to an interface, use the IP address of the interface in the FCIP profile's IP address configuration. Profile numbers range from 1 to 255. The interface associated with a profile can be either of the following:

- Ethernet PortChannel
- Ethernet subinterface slot and port (or slot, port, and VLAN ID)

To modify FCIP profiles and FCIP links on a Gigabit Ethernet interface, follow these steps.

- 
- Step 1** Verify that you are connected to a switch that contains an IPS module.
- Step 2** From Fabric Manager, choose **Switches > ISLs > FCIP** in the Physical Attributes pane. From Device Manager, choose **FCIP** from the IP menu.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 3** Click the **Profiles** tab if it is not already selected. You see the FCIP Profiles dialog box. Any profiles already bound are listed in the table along with their IP addresses.
  - Step 4** Optionally, click the **Create Row** button on Fabric Manager or the **Create** button on Device Manager to add a new profile.
  - Step 5** Enter the profile ID in the ProfileId field.
  - Step 6** Enter the IP address of the interface to which you want to bind the profile.
  - Step 7** Modify the optional TCP parameters, if desired. Refer to Fabric Manager Online Help for explanations of these fields
  - Step 8** Optionally, click the **Tunnels** tab and modify the remote IP address in the Remote IPAddress field for the endpoint to which you want to link.
  - Step 9** Enter the optional parameters, if desired. See the [“Configuring Advanced FCIP Interfaces” section on page 19-13](#).
  - Step 10** Click **Apply Changes** icon to save these changes or Click **Undo Changes** to discard any unsaved changes.
- 

## Verifying Interfaces and Extended Link Protocol

To verify the FCIP interfaces and Extended Link Protocol (ELP) on Device Manager, follow these steps:

- Step 1** Be sure you are connected to a switch that contains an IPS module.
  - Step 2** Select **FCIP** from the Interface menu.
  - Step 3** Click the **Interfaces** tab if it is not already selected. You see the FCIP Interfaces dialog box.
  - Step 4** Click the **ELP** tab if it is not already selected. You see the FCIP ELP dialog box.
- 

## Checking Trunk Status

To check the trunk status for the FCIP interface on Device Manager, follow these steps:

- Step 1** Be sure you are connected to a switch that contains an IPS module.
  - Step 2** Select **FCIP** from the IP menu.
  - Step 3** Click the **Trunk Config** tab if it is not already selected. You see the FCIP Trunk Config dialog box. This shows the status of the interface.
  - Step 4** Click the **Trunk Failures** tab if it is not already selected. You see the FCIP Trunk Failures dialog box.
-



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Modifying FCIP Write Acceleration or FCIP Compression

To modify FCIP Write Acceleration or FCIP compression using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>ISLs &gt; FCIP</b> from the Physical Attributes pane on Fabric Manager. You see the FCIP profiles and links in the Information pane.<br>On Device manager, choose <b>IP &gt; FCIP</b> . You see the FCIP dialog box. |
| <b>Step 2</b> | Click the <b>Tunnels</b> tab. You see the FCIP link information.   |
| <b>Step 3</b> | Check or uncheck the <b>WriteAccelerator</b> check box.  |
| <b>Step 4</b> | Click the <b>IP Compression radio</b> button for the appropriate compression ratio in the dialog box.  |
| <b>Step 5</b> | Click <b>Apply Changes</b> icon to save these changes or click <b>Undo Changes</b> to discard any unsaved changes.   |
- 

## FCIP Tape Acceleration

Tapes are storage devices that store and retrieve user data sequentially. Applications that access tape drives normally have only one SCSI write operation outstanding to it. This single command process limits the benefit of the write acceleration feature when using an FCIP tunnel over a long-distance WAN link. It impacts backup and archive performance because each SCSI write operation does not complete until the host receives a good status response from the tape drive.

The FCIP tape acceleration feature introduced in Cisco MDS SAN-OS Release 2.0(1b) solves this problem. It improves tape backup and archive operations by allowing faster data streaming from the host to the tape over the WAN link.

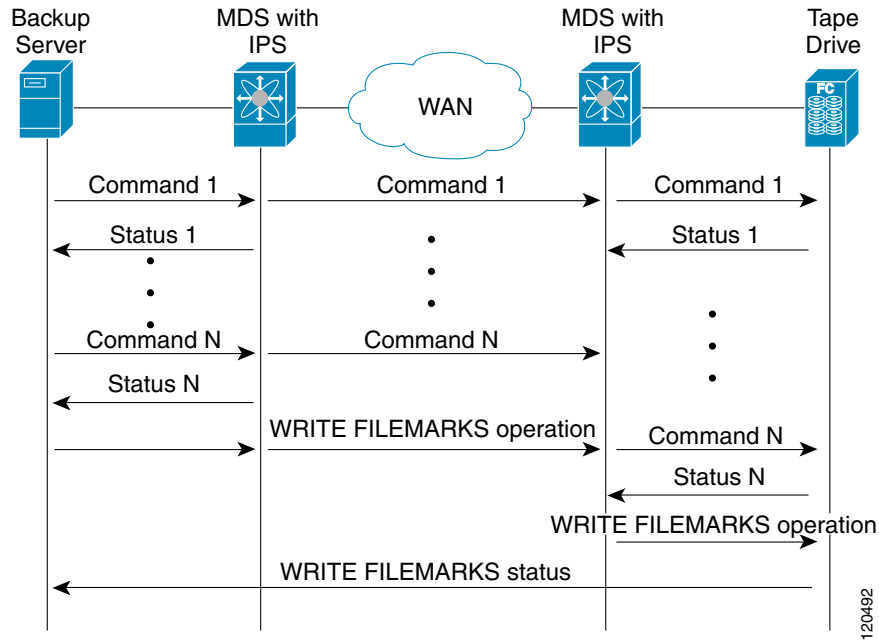
The backup server in [Figure 19-9](#) issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch quickly returns a Transfer Ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch responds to signal the successful completion of the SCSI write operation. This response allows the host to start the next SCSI write operation. This proxy method results in more data being sent over the FCIP tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the use of WAN links.

At the other end of the FCIP tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a Transfer Ready from the tape drive before forwarding the data.

The Cisco MDS SAN-OS provides reliable data delivery to the remote tape drives using tcp/ip over the wan. it maintains write data integrity by allowing the Write Filemarks operation to complete end-to-end without proxying. The write filemarks operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco MDS SAN-OS software.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 19-9 FCIP Link Tape Acceleration**



**Note**

The tape acceleration feature is disabled by default and must be enabled on both sides of the FCIP link. If it is only enabled on one side of the FCIP tunnel, the tunnel is not initialized.



**Tip**

FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths with equal weight between the initiator and the target port. Such a configuration might cause either SCSI discovery failure or broken write or read operations.



**Caution**

When tape acceleration is enabled in an FCIP interface, a FICON VSAN cannot be enabled in that interface. Likewise, if a FCIP interface is up in a FICON VSAN, write acceleration cannot be enabled on that interface.

When you enable the tape acceleration feature, you are automatically enabling the write acceleration feature. When you enable tape acceleration for an FCIP tunnel, the tunnel is reinitialized.

The flow control buffer size specifies the maximum amount of write data that an MDS switch buffers for an FCIP tunnel before it stops the tape acceleration proxying process. The default buffer size is 256 KB and the maximum buffer size is 32 MB.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Enabling FCIP Tape Acceleration

To enable FCIP tape acceleration using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | From Fabric Manager, choose <b>ISLs &gt; FCIP</b> from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.<br><br>From Device Manager, choose <b>IP &gt; FCIP</b> . You see the FCIP dialog box. |
| <b>Step 2</b> | Click the <b>Tunnels</b> tab. You see the FCIP link information.  |
| <b>Step 3</b> | Click the <b>Create Row</b> icon in Fabric Manager or the <b>Create</b> button in Device Manager. You see the FCIP Tunnels dialog box.  |
| <b>Step 4</b> | Set the profile ID in the ProfileID field and the tunnel ID in the TunnelID fields.   |
| <b>Step 5</b> | Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.   |
| <b>Step 6</b> | Check the <b>TapeAccelerator</b> check box.   |
| <b>Step 7</b> | Optionally set the other fields in this dialog box and click <b>Create</b> to create this FCIP link.  |
- 

## Configuring Advanced FCIP Interfaces

You can establish connection to a peer by configuring one or more of the following options for the FCIP interface.

- [Configuring Peers, page 19-13](#)
- [Using B Port Interoperability Mode, page 19-15](#)
- [Configuring E Ports, page 19-18](#)

To establish a peer connection, you must first create the interface. See the “[FCIP Interfaces](#)” section on [page 19-9](#).

## Configuring Peers

To establish a FCIP link with the peer, use one of two options:

- Peer IP address—Configures both ends of the FCIP link. Optionally, you can also use the peer TCP port along with the IP address.
- Special frames—Configures one end of the FCIP link when security gateways are present in the IP network. Optionally, you can also use the port and profile ID along with the IP address.

### Peer IP Address

The basic FCIP configuration uses the peer’s IP address to configure the peer information. You can also specify the peer’s port number to configure the peer information. If you do not specify a port, the default 3225 port number is used to establish connection.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You can configure the required mode for initiating an IP connection. By default, active mode is enabled to actively attempt an IP connection. If you enable the passive mode, the switch does not initiate a TCP connection and merely waits for the peer to connect to it. Ensure that both ends of the FCIP link are not configured as passive mode. If both ends are configured as passive, the connection is not initiated.

You can specify the number of TCP connections from a FCIP link. By default, the switch tries two TCP connections for each FCIP link. You can configure one or two TCP connections. For example, the Cisco PA-FC-1G Fibre Channel port adapter, which has only one TCP connection, interoperates with any switch in the Cisco MDS 9000 Family. One TCP connection is within the specified limit. If the peer initiates one TCP connection, and your MDS switch is configured for two TCP connections, then the software handles it gracefully and moves on with just one connection.

You can instruct the switch to discard packets that are outside the specified time for this FCIP link. By default, this option is disabled in all switches in the Cisco MDS 9000 Family. This option specifies the time range within which packets can be accepted. If the packet arrived within the range specified by this option, the packet is accepted. Otherwise, it is dropped. By default if a packet arrives within a 1000 millisecond interval (+ or -1000 ms), that packet is accepted. Use the **time-stamp** option to enable or disable FCIP timestamps on a packet.



#### Note

If the **time-stamp** option is enabled, be sure to configure NTP on both switches. Refer to the *Cisco MDS 9000 Family Configuration Guide* for information on NTP.

To assign the peer information based on the IP address, port number, or profile ID, follow these steps:

- Step 1** From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.  
From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.
- Step 2** Click the **Tunnels** tab. You see the FCIP link information.
- Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.
- Step 4** Set the ProfileID and TunnelID fields.
- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
- Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
- Step 7** Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.
- Step 8** Optionally, check the **Enable** check box in the Time Stamp section and set the Tolerance field.
- Step 9** Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.

## Special Frames

You can alternatively establish an FCIP link with a peer using an optional protocol called special frames. When special frames are enabled, the peer IP address (and optionally the port or the profile ID) only needs to be configured on one end of the link. Once the connection is established, a special frame is exchanged to discover and authenticate the link.

By default, the special frame feature is disabled.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

Refer to the Fibre Channel IP standards for further information on special frames.

**Tip**

Special frame negotiation provides an additional authentication security mechanism because the link validates the WWN of the peer switch.

To configure special frames, follow these steps:

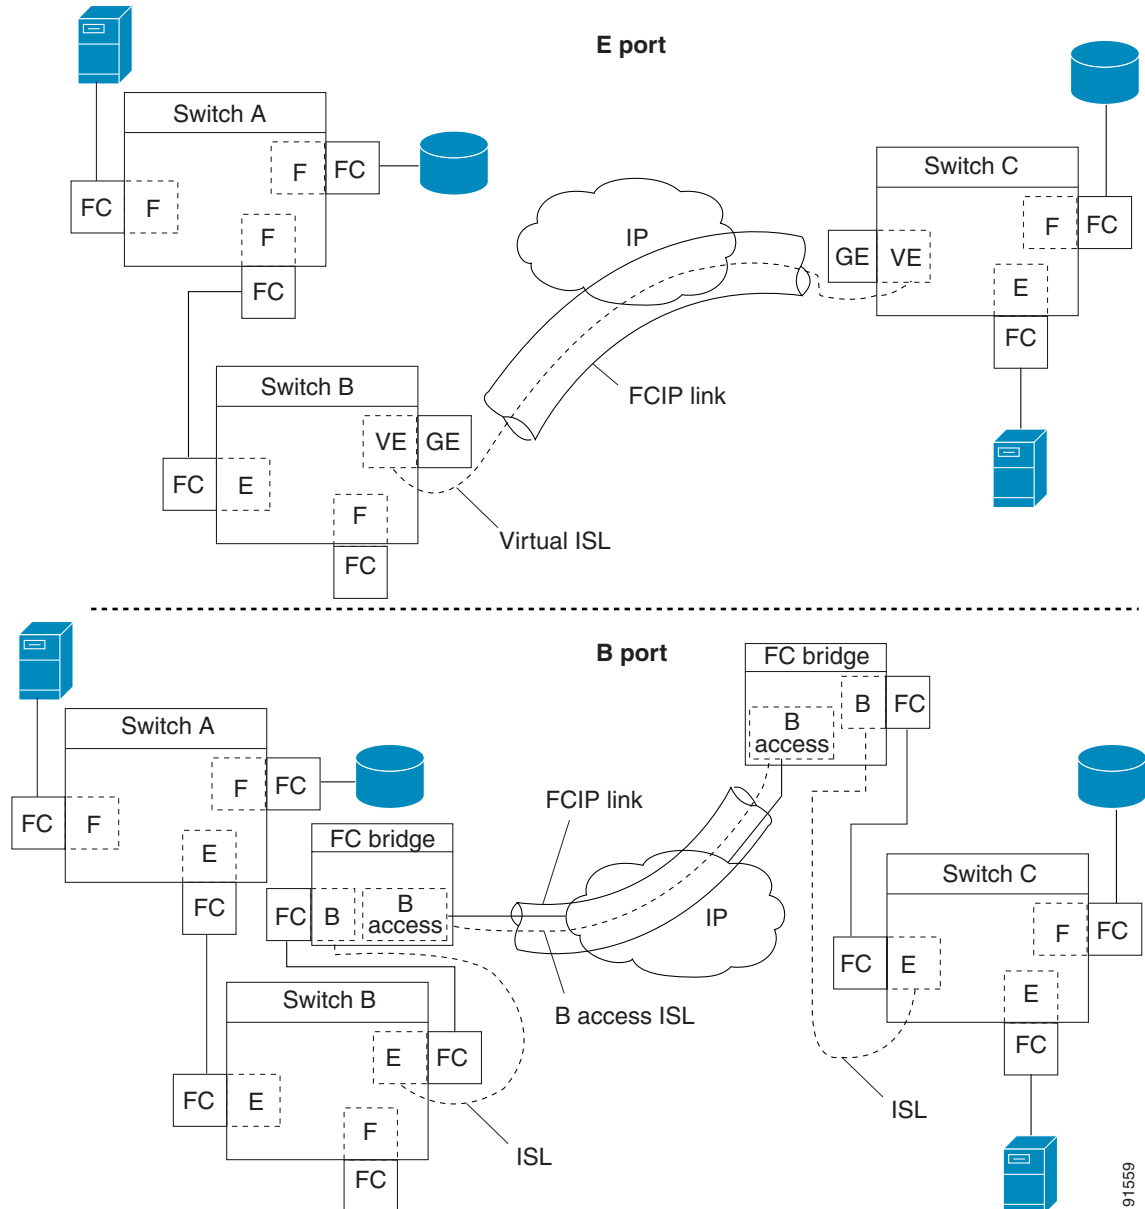
- 
- Step 1** From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.
- From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.
- Step 2** Click the **Tunnels** tab. You see the FCIP link information.
- Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.
- Step 4** Set the ProfileID and TunnelID fields.
- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
- Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
- Step 7** Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.
- Step 8** Check the **Enable** check box in the Special Frames section of the dialog box and set the RemoteWWN and the RemoteProfileID fields.
- Step 9** Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.
- 

## Using B Port Interoperability Mode

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as Cisco's PA-FC-1G Fibre Channel port adapter and the SN 5428-2 storage router, implement a bridge port model to connect geographically dispersed fabrics. This model uses B port as described in the T11 Standard FC-BB-2. [Figure 19-10](#) depicts a typical SAN extension over an IP network.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 19-10 FCIP B Port and Fibre Channel E Port**



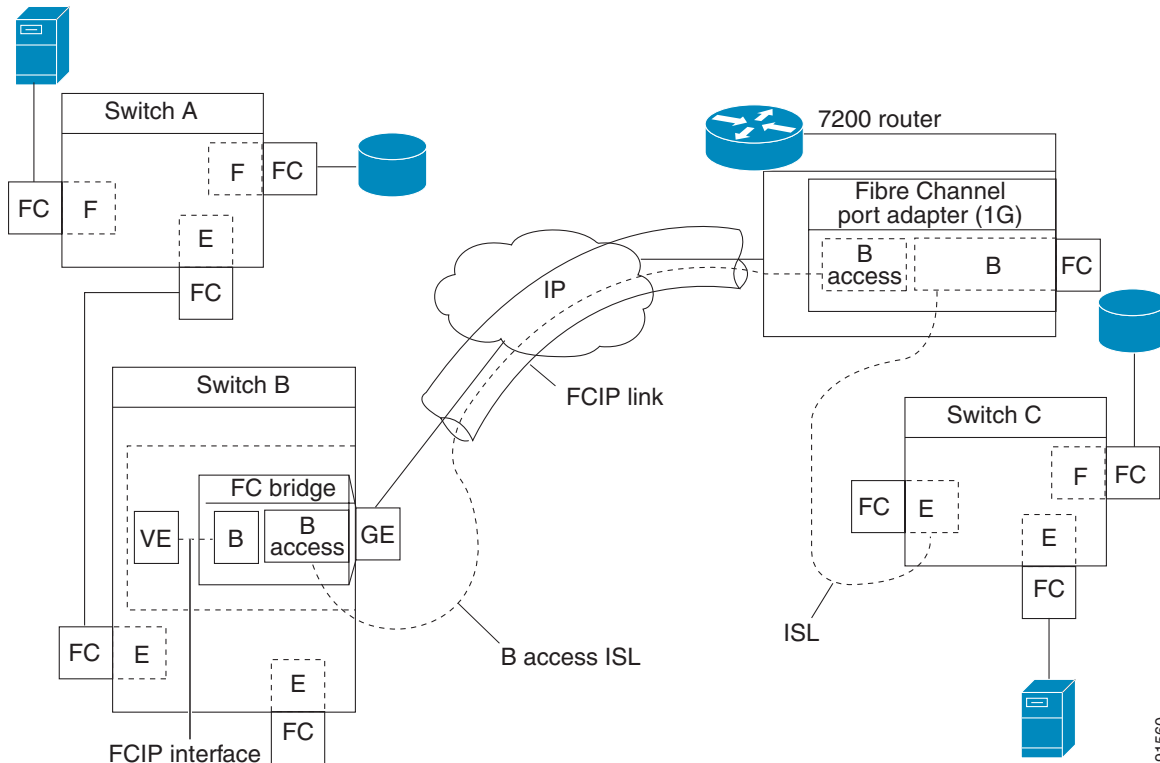
B ports bridge Fibre Channel traffic from one E port to a remote E port without participating in fabric-related activities such as principal switch election, domain ID assignment, and Fibre Channel routing (FSPF). For example, Class F traffic entering a SAN extender does not interact with the B port. The traffic is transparently propagated (bridged) over a WAN interface before exiting the remote B port. This bridge results in both E ports exchanging Class F information that ultimately leads to normal ISL behavior such as fabric merging and routing.

FCIP links between B port SAN extenders do not exchange the same information as FCIP links between E ports, and are therefore incompatible. This is reflected by the terminology used in FC-BB-2: *while VE ports establish a virtual ISL over a FCIP link, B ports use a B access ISL.*

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

The IPS module supports FCIP links that originate from a B port SAN extender device by implementing the B access ISL protocol on a Gigabit Ethernet interface. Internally, the corresponding virtual B port connects to a virtual E port that completes the end-to-end E port connectivity requirement (see Figure 19-11).

**Figure 19-11 FCIP Link Terminating in a B Port Mode**



The B port feature in the IPS module allows remote B port SAN extenders to communicate directly with a Cisco MDS 9000 Family switch, therefore eliminating the need for local bridge devices.

## Configuring B Ports

When a FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled.

To configure B ports for FCIP links using Fabric Manager, follow these steps:

- Step 1** From Fabric Manager, choose **ISLs > FCIP** from the Physical Attributes pane. You see the FCIP profiles and links in the Information pane.  
From Device manager, choose **IP > FCIP**. You see the FCIP dialog box.
- Step 2** Click the **Tunnels** tab. You see the FCIP link information.
- Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the FCIP Tunnels dialog box.
- Step 4** Set the ProfileID and TunnelID fields.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Set the RemoteIPAddress and RemoteTCPPort fields for the peer IP address you are configuring.
  - Step 6** Check the **PassiveMode** check box if you do not want this end of the link to initiate a TCP connection.
  - Step 7** Optionally set the NumTCPCon field to the number of TCP connections from this FCIP link.
  - Step 8** Check the **Enable** check box in the B Port section of the dialog box and optionally check the **KeepAlive** check box if you want a response sent to an ELS Echo frame received from the FCIP peer.
  - Step 9** Optionally set the other fields in this dialog box and click **Create** to create this FCIP link.
- 

## Configuring E Ports

All configuration commands that apply to E ports, also apply to FCIP interfaces. The following features are also available FCIP interfaces:

- VSANs—FCIP interfaces can be a member of any VSAN.
- Trunk mode and trunk allowed VSANs can be configured.
- PortChannels.
  - Multiple FCIP links can be bundled into a Fibre Channel PortChannel.
  - FCIP links and Fibre Channel links cannot be combined in one PortChannel.
- FSPF.
- Fibre Channel domains (fcdomains).
- Zone merge—The zone database can be imported or exported from the adjacent switch.

## FCIP High Availability

The following high availability solutions are available for FCIP configurations:

- [Fibre Channel PortChannels, page 19-19](#)
- [FSPF, page 19-19](#)
- [VRRP, page 19-20](#)
- [Ethernet PortChannels, page 19-20](#)
- [Ethernet PortChannels and Fibre Channel PortChannels, page 19-21](#)




**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Fibre Channel PortChannels

Figure 19-12 provides an example of a PortChannel-based load balancing configuration. To perform this configuration, you need two IP addresses on each SAN island. This solution addresses link failures.

**Figure 19-12 PortChannel Based Load Balancing**



90857


The following characteristics set Fibre Channel PortChannel solutions apart from other solutions:

- The entire bundle is one logical (E)ISL link.
- All FCIP links in the PortChannel should be across the same two switches.
- The Fibre Channel traffic is load balanced across the FCIP links in the PortChannel.

## FSPF

Figure 19-13 displays a FSPF-based load balancing configuration example. This configuration requires two IP addresses on each SAN island, and addresses IP and FCIP link failures.

**Figure 19-13 FSPF-Based Load Balancing**



90858

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**


The following characteristics set FSPF solutions apart from other solutions:

- Each FCIP link is a separate (E)ISL.
- The FCIP links can connect to different switches across two SAN islands.
- The Fibre Channel traffic is load balanced across the FCIP link.

## VRRP

Figure 19-14 displays a VRRP-based high availability FCIP configuration example. This configuration requires at least two physical Gigabit Ethernet ports connected to the Ethernet switch on the island where you need to implement high availability using VRRP.

**Figure 19-14 VRRP-Based High Availability**




The following characteristics set VRRP solutions apart from other solutions:

- If the active VRRP port fails, the standby VRRP port takes over the VRRP IP address.
- When the VRRP switchover happens, the FCIP link automatically disconnects and reconnects.
- This configuration has only one FCIP (E)ISL link.

## Ethernet PortChannels

Figure 19-15 displays an Ethernet PortChannel-based high availability FCIP example. This solution addresses the problem caused by individual Gigabit Ethernet link failures.

**Figure 19-15 Ethernet PortChannel-Based High Availability**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***


The following characteristics set Ethernet PortChannel solutions apart from other solutions:

- The Gigabit Ethernet link level redundancy ensures a transparent failover if one of the Gigabit Ethernet links fails.
- Two Gigabit Ethernet ports in one Ethernet PortChannel appear like one logical Gigabit Ethernet link.
- The FCIP link stays up during the failover.

## Ethernet PortChannels and Fibre Channel PortChannels

Ethernet PortChannels offer Ethernet-level redundancy and Fibre Channel PortChannels offer (E)ISL-level redundancy. FCIP is unaware of any Ethernet PortChannels or Fibre Channel PortChannels. Fibre Channel PortChannels are unaware of any Ethernet PortChannels, and there is no mapping between the two (see [Figure 19-16](#)).

**Figure 19-16** *PortChannels at the Fibre Channel and Ethernet Levels*



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## iSCSI Configuration

---

Cisco MDS 9000 Family IP storage services (IPS) modules extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. The switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol. The IPS modules include the IPS-8, IPS-4, and the MPS-14/2 modules.



**Note**

iSCSI features are specific to the IPS-8 modules running Cisco MDS SAN-OS Release 1.1(x) or later, the IPS-4 modules running Cisco MDS SAN-OS Release 1.3(4a), and the Gigabit Ethernet ports on the MPS-14/2 module running Cisco MDS SAN-OS Release 2.0(x) or later.

This chapter includes the following sections:

- [Configuring iSCSI, page 20-1](#)
- [Configuring Storage Name Services, page 20-23](#)

## Configuring iSCSI

This section includes the following topics:


- [About iSCSI, page 20-1](#)
- [Enabling iSCSI, page 20-5](#)
- [Using the iSCSI Wizard, page 20-5](#)
- [Presenting iSCSI Hosts as Virtual Fibre Channel Hosts, page 20-11](#)
- [Creating a Statically Mapped iSCSI Initiator, page 20-13](#)
- [iSCSI Proxy Initiators, page 20-14](#)
- [Access Control in iSCSI, page 20-16](#)
- [iSCSI User Authentication, page 20-17](#)
- [Advanced iSCSI Configuration, page 20-19](#)

## About iSCSI

The IPS module provides transparent SCSI routing by default. IP hosts using the iSCSI protocol can transparently access targets on the Fibre Channel network. [Figure 20-1](#) provides an example of a typical configuration of iSCSI hosts with access to a Fibre Channel SAN.


**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 20-1 Typical IP to Fibre Channel SAN Configuration**



IPS modules enable you to create virtual iSCSI targets and then map them to physical Fibre Channel targets available in the Fibre Channel SAN. They present the Fibre Channel targets to IP hosts as if the physical targets were attached to the IP network (see [Figure 20-2](#)).


**Figure 20-2 iSCSI View**



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

In conjunction with presenting Fibre Channel targets to iSCSI hosts, the IPS module presents each iSCSI host as a Fibre Channel host (in transparent mode), that is, a host bus adapter (HBA) to the Fibre Channel storage device. The storage device responds to each IP host as if it were a Fibre Channel host connected to the Fibre Channel network (see Figure 20-3).

**Figure 20-3 Fibre Channel SAN View**



**Note**

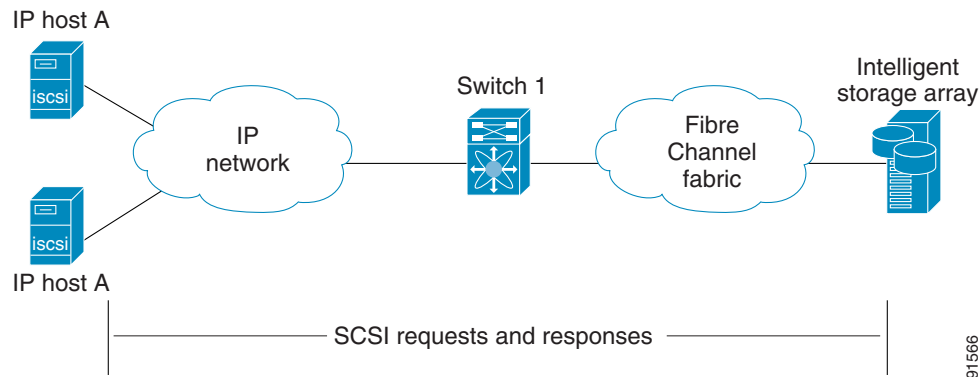
Refer to the IETF standards for IP storage at <http://www.ietf.org> for information on the iSCSI protocol

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Routing iSCSI Requests and Responses

The iSCSI feature consists of routing iSCSI requests and responses between hosts in an IP network and Fibre Channel storage devices in the Fibre Channel SAN that are accessible from any Fibre Channel interface of the Cisco MDS 9000 Family switch (see [Figure 20-4](#)).

**Figure 20-4** Routing iSCSI Requests and Responses for Transparent iSCSI Routing



Each iSCSI host that requires access to storage through the IPS module needs to have a compatible iSCSI driver installed. (The Cisco.com website at <http://www.cisco.com/kobayashi/sw-center/sw-stornet.shtml> provides a list of compatible drivers). Using the iSCSI protocol, the iSCSI driver allows an iSCSI host to transport SCSI requests and responses over an IP network. From the host operating system perspective, the iSCSI driver appears to be a SCSI transport driver similar to a Fibre Channel driver for a peripheral channel in the host. From the storage device perspective, each IP host appears as a Fibre Channel host.

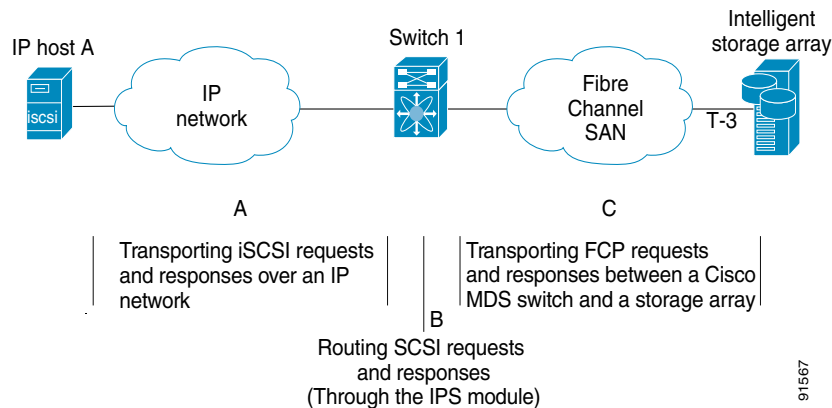
Routing SCSI from the IP host to the Fibre Channel storage device consists of the following main actions (see [Figure 20-5](#)):

- The iSCSI requests and responses are transported over an IP network between the hosts and the IPS module.
- The SCSI requests and responses are routed between the hosts on an IP network and the Fibre Channel storage device (converting iSCSI to FCP and vice versa). The IPS module performs this routing.
- The FCP requests or responses are transported between the IPS module and the Fibre Channel storage devices.



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 20-5**      **Transparent SCSI Routing Actions**



**Note**

FCP (the Fibre Channel equivalent of iSCSI) carries SCSI commands over a Fibre Channel SAN.

## Enabling iSCSI

To begin configuring the iSCSI feature, you must explicitly enable iSCSI on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the iSCSI feature are only available when iSCSI is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable iSCSI on a switch using Fabric Manager, follow these steps:

- Step 1** Choose **End Devices > iSCSI** from the Physical Attributes pane. You see the iSCSI tables in the Information pane.
- Step 2** Click the **Control** tab if it is not already displayed. You see the iSCSI enable status for all switches in the fabric that contain IPS ports.
- Step 3** Choose **enable** from the Command column for each switch that you want to enable iSCSI on.
- Step 4** Click the **Apply Changes** icon to save these changes or click the **Undo Changes** icon to remove all changes without saving them.

## Using the iSCSI Wizard

To use the iSCSI wizard in Fabric Manager, follow these steps:

- Step 1** Choose the **iSCSI Setup Wizard** icon.
- Step 2** Select an existing iSCSI initiator or add the iSCSI node name or IP address for a new iSCSI initiator.
- Step 3** Select the switch for this iSCSI initiator if you are adding a new iSCSI initiator and click **Next**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Select the VSAN and targets to associate with this iSCSI initiator, as shown in [Figure 20-6](#) and click **Next**.

**Figure 20-6**      **Select Targets**

iSCSI Wizard - /SAN 1/Fabric 172.22.31.190

2 of 3: Select Targets

Select targets to be associated with iSCSI initiator

VSAN: VSAN0001

Available

Name	Switch Interface	Fcid	iSCSI Name

▼ Add    ▲ Remove

Selected

Name	Switch Interface	Fcid	iSCSI Name
Seagate 21:00:00:20:37:6f:db:63	172.22.31.184 fc4/31	0x6c0101	
Seagate 21:00:00:04:cf:fb:42:f8	172.22.31.184 fc4/30	0x6c0001	

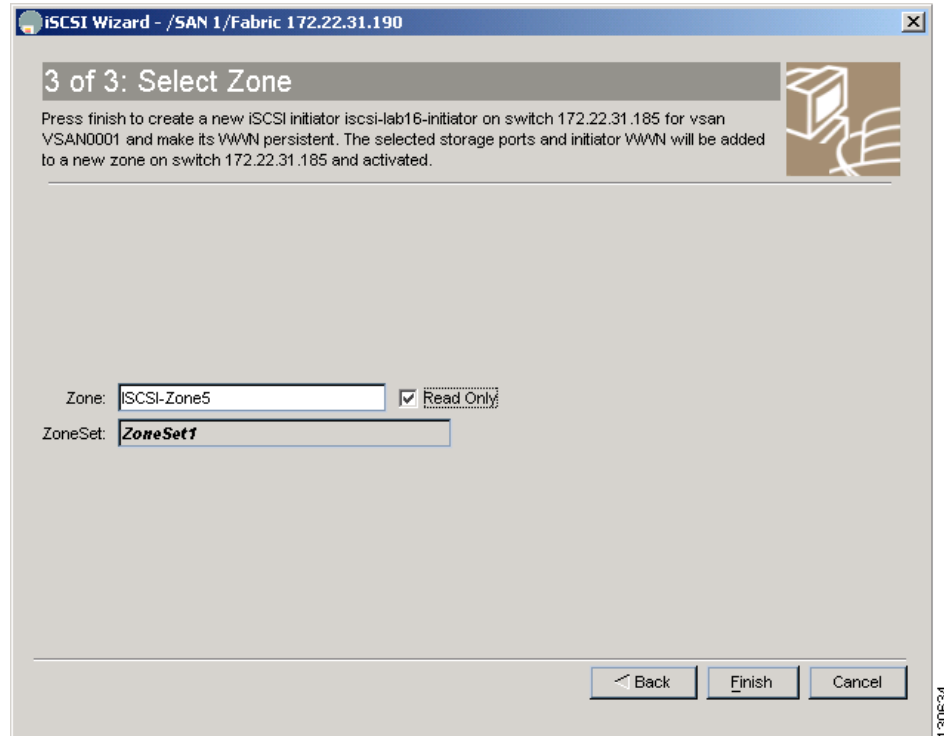
Back    Next    Cancel

130633

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Set the zone name for this new iSCSI zone and optionally check the **Read Only** check box, as shown in Figure 20-7.

**Figure 20-7 Select Zone**



- Step 6** Click **Finish** to create this iSCSI initiator or click **Cancel** to close the wizard without creating the iSCSI initiator. If created, the target VSAN is added to the iSCSI host VSAN list.

## Presenting Fibre Channel Targets as iSCSI Targets

The IPS module presents physical Fibre Channel targets as iSCSI targets allowing them to be accessed by iSCSI hosts. It does this in one of two ways:

- Dynamic importing—use if all logical units (LUs) in all Fibre Channel storage targets are made available to iSCSI hosts (subject to VSAN and zoning).
- Static importing—use if iSCSI hosts are restricted to subsets of LUs in the Fibre Channel targets and additional iSCSI access control is needed (see the [“Access Control in iSCSI” section on page 20-16](#)). Also, static importing allows automatic failover if the LUs of the Fibre Channel targets are reached by redundant Fibre Channel ports (see the [“High Availability Static Target Importing” section on page 20-10](#)).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



**Note**

The IPS module does not import Fibre Channel targets to iSCSI by default. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have a configured name. Targets that are not statically imported are advertised with the name created by the conventions explained in this section.

## Dynamically Importing Fibre Channel Targets

The IPS module maps each physical Fibre Channel target port as one iSCSI target. That is, all LUs accessible through the physical storage target port are available as iSCSI LUs with the same LU number (LUN) as in the storage target.

For example, if an iSCSI target was created for a Fibre Channel target port with pWWN 31:00:11:22:33:44:55:66 and that pWWN contains LUN 0 through 2, those LUNs would become available to an IP host as LUNs 0 through 2 as well.



**Note**

If you have configured a switch name, then the switch name is used instead of the management IP address. If you have not configured a switch name, the management IP address is used.

The iSCSI target node name is created automatically using the iSCSI qualified name (IQN) format. The iSCSI qualified name is restricted to a maximum name length of 223 alphanumeric characters and a minimum length of 16 characters.

The IPS module creates an IQN formatted iSCSI node name using the following conventions:

- IPS ports that are not part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.<mgmt-ip-address>.<slot#>-<port#>-<sub-intf#>.<Target-pWWN>
```

- IPS ports that are part of a VRRP group use this format:

```
iqn.1987-05.com.cisco:05.vrrp-<vrrp-ID#>-<vrrp-IP-addr>.<Target-pWWN>
```

- Ports that are part of a PortChannel use this format:

```
iqn.1987-02.com.cisco:05.<mgmt-ip-address>.pc-<port-ch-sub-intf#>.<Target-pWWN>
```



**Note**

With this format, each IPS port in a Cisco MDS 9000 Family switch creates a different iSCSI target node name for the same Fibre Channel target.

## Configuring Dynamic Importing with Device Manager

To dynamically import Fibre Channel targets as iSCSI targets, follow these steps:

- Step 1** Choose **IP > iSCSI** in Device Manager. You see the iSCSI dialog box.
- Step 2** Choose the **Targets** tab to display a list of existing iSCSI targets.
- Step 3** Check the **Dynamically Import FC Targets** check box.
- Step 4** Click **Apply** to save this change or click **Cancel** to close the dialog box without saving any changes.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Creating a Static iSCSI Virtual Target

You can manually (statically) create an iSCSI target and assign a node name to it. A statically mapped iSCSI target can either contain the whole FC target port, or it can contain one or more LUs from a Fibre Channel target port.

You can limit the Gigabit Ethernet interfaces over which static iSCSI targets are advertised. By default iSCSI targets are advertised on all Gigabit Ethernet interfaces, subinterfaces, PortChannel interfaces, and PortChannel subinterfaces.

To create a static iSCSI virtual target for the Fibre Channel target port, follow these steps

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>IP &gt; iSCSI</b> in Device Manager. You see the iSCSI dialog box.   |
| <b>Step 2</b> | Choose the <b>Targets</b> tab to display a list of existing iSCSI targets.   |
| <b>Step 3</b> | Click <b>Create</b> to create an iSCSI target. You see the Create iSCSI Targets dialog box.  |
| <b>Step 4</b> | Set the iSCSI target node name in the iSCSI Name field, in IQN format.   |
| <b>Step 5</b> | Set the Port WWN field for the Fibre Channel target port you are mapping.  |
| <b>Step 6</b> | Choose the <b>List</b> radio button and set the iSCSI initiator node names or IP Addresses that you want this virtual iSCSI target to access, or choose the <b>All</b> radio button to let the iSCSI target access all iSCSI initiators. See the <a href="#">“Access Control in iSCSI” section on page 20-16</a> . |
| <b>Step 7</b> | Choose the <b>Selected from List</b> radio button and check each interface you want to advertise the iSCSI targets on or choose the <b>All</b> radio button to advertise all interfaces.   |
| <b>Step 8</b> | Click <b>Apply</b> to save this change or click <b>Cancel</b> to close the dialog box without saving any changes.  |
- 

See the [“iSCSI-Based Access Control” section on page 20-17](#) for more information on controlling access to statically imported targets.


For multiple interfaces configured with iSNS (see the [“Configuring Storage Name Services” section on page 20-23](#)), a different static virtual target name has to be created for each interface tagged to an iSNS profile and each static virtual target must be advertised only from one interface (see the [“Configuring Storage Name Services” section on page 20-23](#))

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## High Availability Static Target Importing

Statically imported iSCSI targets have an additional option to provide a secondary pWWN for the Fibre Channel target. This can be used when the physical Fibre Channel target is configured to have an LU visible across redundant ports. When the active port fails, the secondary port becomes active and the iSCSI session switches to use the new active port (see Figure 20-8).

**Figure 20-8 Static Target Importing Through Two Fibre Channel Ports**



In Figure 20-8, you can create a virtual iSCSI target that is mapped to both pWWN1 and pWWN2 to provide redundant access to the Fibre Channel targets.

The failover to secondary port is done transparently by the IPS port without impacting the iSCSI session from the host. All outstanding I/O are terminated with a check condition status when the primary port fails. New I/O received while the failover has not completed will receive a busy status.



**Tip**

If you use LUN mapping, you can define a different secondary Fibre Channel LUN if the LU number is different.

Refer to the *Cisco MDS 9000 Family Configuration Guide* for details on setting the secondary pWWN.

Enable the **revert to primary port** option to direct the IPS port to switch back to the primary port when the primary port is up again. If this option is disabled (default) and the primary port is up again after a switchover, the old sessions will remain with the secondary port and does not switch back to the primary port. However, any new session will use the primary port. This is the only situation when both the primary and secondary ports are used at the same time.

To enable the revert to primary port option, follow these steps:

- Step 1** From Fabric Manager, choose **End Devices > iSCSI** from the Physical Attributes pane. You see the iSCSI tables in the Information pane.  
From Device Manager, choose **IP > iSCSI**. You see the iSCSI dialog box.
- Step 2** Choose the **Targets** tab to display a list of existing iSCSI targets.
- Step 3** Check the **RevertToPrimaryPort** check box to enable this option.
- Step 4** Set the iSCSI target node name in the iSCSI Name field, in IQN format.
- Step 5** Click the **Apply Changes** icon in Fabric Manager or the **Apply** button in Device Manager to save this change or click **Cancel** to close the dialog box without saving any changes.


**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Configuring the Trespass Feature

In addition to the high availability of statically imported iSCSI targets, the trespass feature is available (as of Cisco MDS SAN-OS Release 1.3(x)) to enable the export of LUs, on an active port failure, from the active to the passive port of a statically imported iSCSI target.

In physical Fibre Channel targets, which are configured to have LUs visible over two Fibre Channel N-ports, when the active port fails, the passive port takes over. Some physical Fibre Channel targets require that the **trespass** command be issued to export the LUs from the active port to the passive port. A statically imported iSCSI target's secondary pWWN option and an additional option of enabling the trespass feature is available for a physical Fibre Channel target with redundant ports. When the active port fails, the passive port becomes active, and if the trespass feature is enabled, the Cisco MDS switch issues a **trespass** command to the target to export the LUs on the new active port. The iSCSI session switches to use the new active port and the exported LUs are accessed over the new active port (see Figure 20-9).

**Figure 20-9 Virtual Target with an Active Primary Port**



To configure the trespass feature, follow these steps:

- 
- Step 1** From Fabric Manager, choose **End Devices > iSCSI** from the Physical Attributes pane. You see the iSCSI tables in the Information pane.  
From Device Manager, choose **IP > iSCSI**. You see the iSCSI dialog box.
  - Step 2** Choose the **Targets** tab to display a list of existing iSCSI targets.
  - Step 3** Check the **Trespass Mode** check box to enable this option.
  - Step 4** Click the **Apply Changes** icon in Fabric Manager or the **Apply** button in Device Manager to save this change or click **Cancel** to close the dialog box without saving any changes.
- 

## Presenting iSCSI Hosts as Virtual Fibre Channel Hosts

The iSCSI hosts are mapped to virtual Fibre Channel hosts in one of two ways (see Figure 20-3):

- Dynamic mapping (default)—use if no access control is done on the Fibre Channel target. An iSCSI host may use different pWWNs each time it connects to a Fibre Channel target.
- Static mapping—use if an iSCSI host should always have the same pWWN or nWWN each time it connects to a Fibre Channel target.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Dynamic Mapping

When an iSCSI host connects to the IPS module using the iSCSI protocol, a virtual N port is created for the host. The nWWNs and pWWNs are dynamically allocated from the switch's Fibre Channel WWN pool. The IPS module registers this N port in the Fibre Channel SAN. The IPS module continues using that nWWN and pWWN to represent this iSCSI host until it no longer has a connection to any iSCSI target through that IP storage port.

At that point, the virtual Fibre Channel host is taken offline from the Fibre Channel SAN and the nWWNs and pWWNs are released back to the switch's Fibre Channel WWN pool. These addresses become available for assignment to other iSCSI hosts requiring access to Fibre Channel SANs. When a dynamically mapped iSCSI initiator has multiple sessions to multiple Fibre Channel targets, each session can use the same pWWN and nWWN as long as it uses the same node name in the iSCSI login message.

### Initiator Identification

By default, the switch uses the iSCSI node name to identify the initiator.

An iSCSI initiator is identified in one of two ways:

- By iSCSI node name—An initiator with multiple IP addresses (multiple interface cards—NICs or multiple network interfaces) has one virtual N port, assuming it uses the same iSCSI initiator name from all interfaces.
- By IP address—A virtual N port is created for each IP address it uses to log in to iSCSI targets.

## Static Mapping

Use the static mapping method to obtain the same nWWN and pWWNs for the iSCSI host each time it connects to the IPS module.

Static mapping can be used on the IPS module to access intelligent Fibre Channel storage arrays that have access control and LUN mapping or masking configuration based on the initiator's pWWNs and/or nWWNs.



#### Note

If an iSCSI host connects to multiple IPS ports, each port independently creates one virtual N port for the host. If static mapping is used, enough pWWNs should be configured for as many IPS ports to which a host connects.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You can implement static mapping in one of two ways:

- **Manual assignment**—You can specify your own unique WWN by providing them during the configuration process.
- **System assignment**—When a static mapping configuration is created, one nWWN and/or one or more pWWNs are allocated from the switch's Fibre Channel WWN pool and the mapping is kept permanent.



**Tip**

We recommend using the **system assignment** option. If you manually assign a WWN, you must ensure its uniqueness (see the “[Configuring World Wide Names](#)” section on page 23-3).



**Note**

If a system-assign option is used to configure WWNs for an iSCSI initiator, when the configuration is backed up to an ASCII file the system-assigned WWNs are also saved. Subsequently if you perform a **write erase**, you must manually delete the WWN configuration from the ASCII file.

## Assigning VSAN Membership to iSCSI Hosts

By default, a host is only in VSAN 1 (default VSAN). You can configure an iSCSI host to be a member of one or more VSANs. The IPS module creates one Fibre Channel virtual N port in each VSAN to which the host belongs.



**Note**

When an initiator is configured in any other VSAN (other than VSAN 1), for example VSAN 2, the initiator is automatically removed from VSAN 1. If you also want it to be present in VSAN 1, you must explicitly configure the initiator in VSAN 1.

All dynamic iSCSI initiators are members of VSAN 1. The port VSAN of an iSCSI interface is the default VSAN for all dynamic iSCSI initiators. All dynamic iSCSI initiators are members of the port VSAN of the iSCSI interface. The default port VSAN of an iSCSI interface is VSAN 1.

To modify the VSANs assigned to an iSCSI interface using Device Manager, follow these steps:

- Step 1** Select **Interfaces > Ethernet or iSCSI**. You see the interfaces dialog box.
- Step 2** Click the **iSCSI** tab. You see the iSCSI interface configuration table.
- Step 3** Double-click the PortVSAN column to modify the default port VSAN.
- Step 4** Click **Apply** to save these changes, or click **Cancel** to discard changes.

## Creating a Statically Mapped iSCSI Initiator

To create a statically mapped iSCSI initiator using Device Manager, follow these steps:

- Step 1** Select **IP > iSCSI**. You see the iSCSI configuration dialog box.
- Step 2** Select **Initiators** tab if it is not already displayed.
- Step 3** Click **Create** to create an iSCSI initiator.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Set the iSCSI node name or IP address and VSAN membership.
  - Step 5** In the Node WWN section, check the **Persistent** check box.
  - Step 6** Check the **System Assigned** check box if you want the switch to assign the nWWN. Or leave this unchecked and set Static WWN field.
  - Step 7** In the Port WWN section, check the **Persistent** check box if you want to statically map pWWNs to the iSCSI initiator.
  - Step 8** If persistent, check the **System Assigned** check box and set the number of pWWNs to reserve for this iSCSI initiator if you want the switch to assign pWWNs.  
Or leave this unchecked and set one or more pWWNs for this iSCSI initiator.
  - Step 9** Optionally set the AuthUser field if authentication is enabled. See the [“iSCSI User Authentication” section on page 20-17](#).
  - Step 10** Click **Create** to create this iSCSI initiator or click **Cancel** to close the dialog box without creating a new iSCSI initiator.
- 

## iSCSI Proxy Initiators



### Note

When an interface is in the proxy initiator mode, you can only configure Fibre Channel access control (zoning) based on the Fibre Channel interface attributes—the WWN pair and available FCIDs. You cannot configure zoning using iSCSI attributes such as the IP address or the iQN name of the iSCSI initiator. To enforce initiator-based access control, use iSCSI based access control (see the [“Access Control in iSCSI” section on page 20-16](#)).


---

By default, each iSCSI initiator appears as one Fibre Channel initiator in transparent mode in the Fibre Channel fabric. For some storage arrays, this appearance requires the initiator's pWWN to be manually configured for access control purposes. This process can be quite cumbersome. The proxy initiator feature allows all iSCSI initiators to connect through one IPS port making it appear as one Fibre Channel port per VSAN. It simplifies the task of configuring the pWWN for each new initiator on the storage array, and of configuring Fibre Channel access control such as zoning. This feature along with static target importing (using LUN mapping) results in the configuration being performed only on the switch

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**


when a new iSCSI host is added. On the storage array, all LUNs that are used by iSCSI initiators are configured to allow access by the proxy initiator's pWWN. From the iSCSI perspective, this configuration is no different from the default mode (see [Figure 20-10](#)).

**Figure 20-10 iSCSI View of a Proxy Initiator**



From the Fibre Channel perspective, only one Fibre Channel initiator is visible per VSAN (see [Figure 20-11](#)).

**Figure 20-11 Fibre Channel View with a Proxy Initiator**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring the iSCSI Proxy Initiator

To configure the proxy initiator, follow these steps:

- 
- Step 1** From Fabric Manager, select **Interfaces > FC Logical** from the Physical Attributes pane. You see the Interface tables in the Information pane.
- From Device Manager, select **Interfaces > Ethernet or iSCSI**. You see the interfaces dialog box.
- Step 2** Click the **iSCSI** tab. You see the iSCSI interface configuration table.
- Step 3** In the Initiator Proxy Mode section, check the **Enable** check box.
- Step 4** Click the **Apply Changes** icon in Fabric Manager or click **Apply** in Device Manager to save these changes. Or click **Undo Changes** in Fabric Manager or click **Cancel** in Device Manager to discard changes.
- 

## Access Control in iSCSI

You can control access to each statically mapped iSCSI target by specifying a list of IPS ports on which it is advertised and specifying a list of iSCSI initiator node names allowed to access it. Fibre Channel zoning-based access control and iSCSI-based access control are the two mechanisms by which access control can be provided for iSCSI. Both methods can be used simultaneously.



### Note

This access control is in addition to the existing Fibre Channel access control. The iSCSI initiator has to be in the same VSAN and zone as the physical Fibre Channel target.

## Fibre Channel Zoning-Based Access Control

Zoning is an access control mechanism within a VSAN. The zoning implementation on the switch extends the VSAN and zoning concepts from the Fibre Channel domain to cover the iSCSI domain. This extension includes both iSCSI and Fibre Channel features and provides a uniform, flexible access control across a SAN. There are two Fibre Channel zoning access control mechanisms--static and dynamic.

- **Static**—Statically map the iSCSI host to Fibre Channel virtual N port(s). This creates permanent nWWNs and pWWNs. Next, configure the assigned pWWN into zones, similar to adding a regular Fibre Channel host pWWN to a zone.
- **Dynamic**—Add the iSCSI host initiator node name as a member of a zone. When the IP host Fibre Channel virtual N port is created and the Fibre Channel address (nWWNs and pWWNs) is assigned, Fibre Channel zoning is enforced.

To register an iSCSI host initiator as a member of a zone using Fabric Manager, follow these steps:

- 
- Step 1** Select **Zone > Edit Local Full Zone Database**.
- Step 2** Select the VSAN and zone you want to add the iSCSI host initiator to.
- Step 3** From the list of available devices including iSCSI host initiators, click the initiators that you want to add to the zone and click the **Add to Zone or Alias** button.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 4** Click **Close** to close the dialog box.

---

## iSCSI-Based Access Control

For static iSCSI targets, you can manually configure a list of iSCSI initiators that are allowed to access the targets. The iSCSI initiator is identified by the iSCSI node name or the IP address of the iSCSI host.

By default, static virtual iSCSI targets are not accessible to any iSCSI host. You must explicitly configure accessibility to allow a virtual iSCSI target to be accessed by all hosts. The initiator access list can contain one or more initiators. Each initiator is identified by one of the following:

- iSCSI node names
- IP addresses
- IP subnets

See the [“Creating a Static iSCSI Virtual Target” section on page 20-9](#) to configure access control using a list of authorized initiators.

## Enforcing Access Control

IPS modules use both iSCSI node name-based and Fibre Channel zoning-based access control lists to enforce access control during iSCSI discovery and iSCSI session creation.

- iSCSI discovery—When an iSCSI host creates an iSCSI discovery session and queries for all iSCSI targets, the IPS module returns only the list of iSCSI targets this iSCSI host is allowed to access based on the access control policies discussed in the previous section.
- iSCSI session creation—When an IP host initiates an iSCSI session, the IPS module verifies if the specified iSCSI target (in the session login request) is a static mapped target, and if true, verifies if the IP host's iSCSI node name is allowed to access the target. If the IP host does not have access, its login is rejected.

The IPS module then creates a Fibre Channel virtual N port (the N port may already exist) for this IP host and does a Fibre Channel name server query for the FCID of the Fibre Channel target pWWN that is being accessed by the IP host. It uses the IP host virtual N port's pWWN as the requester of the name server query. Thus, the name server does a zone-enforced query for the pWWN and responds to the query.

If the FCID is returned by the name server, then the iSCSI session is accepted. Otherwise, the login request is rejected.



### Note

If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts. Refer to the *Cisco MDS 9000 Family Configuration Guide* to configure persistent FC IDs.

## iSCSI User Authentication

The IPS module supports the iSCSI authentication mechanism to authenticate iSCSI hosts that request access to storage. When iSCSI authentication is enabled, the iSCSI hosts must provide user name and password information each time an iSCSI session is established.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

Only the Challenge Handshake Authentication Protocol (CHAP) authentication method is supported.

## No Authentication

If no authentication is configured, local authentication is used.

Set the iSCSI authentication method to **none** to configure a network with no authentication. See the [“Configuring an Authentication Mechanism”](#) section on page 20-18.

## Configuring an Authentication Mechanism

During an iSCSI login, both the iSCSI initiator and target have the option to authenticate each other. By default, the IPS module allows either CHAP authentication or no authentication from iSCSI hosts.

**Note**

The authentication for a Gigabit Ethernet interface or subinterface configuration overrides the authentication for the global interface configuration.

To configure an authentication method for iSCSI, follow these steps:

- Step 1** In Fabric Manager, select **End Devices > iSCSI** from the Physical Attributes pane. You see the iSCSI tables in the Information pane.  
In Device Manager, select **IP > iSCSI**. You see the iSCSI dialog box.
- Step 2** Click the **Global** tab. You see the iSCSI authentication configuration table.
- Step 3** In Fabric Manager, select **chap** or **none** from the authMethod column.  
Or in Device manager, check the **Chap** check box to configure DH-CHAP authentication, or check **none** for no authentication.
- Step 4** Click the **Apply Changes** icon in Fabric Manager or click **Apply** in Device Manager to save these changes, or click **Undo Changes** in Fabric Manager or click **Cancel** in Device Manager to discard changes.

## Configuring an iSCSI RADIUS Server

To configure an iSCSI RADIUS server, follow these steps:

- Step 1** Configure the RADIUS server to allow access from the Cisco MDS switch's management Ethernet IP address.
- Step 2** Configure the shared secret for the RADIUS server to authenticate the Cisco MDS switch.
- Step 3** Configure the iSCSI users and passwords on the RADIUS server.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Advanced iSCSI Configuration

Advanced configuration options are available for iSCSI interfaces on a per-IPS port basis. Cisco MDS switches support the following advanced features for iSCSI interfaces:

- iSCSI listener port—You can configure the TCP port number for the iSCSI interface which listens for new TCP connections. The default port number is 3260. Following that, the iSCSI port only accepts TCP connections on the newly configured port  
See the )
- TCP tuning parameters—You can configure the following TCP parameters.
  - The minimum retransmit timeout, keepalive timeout, maximum retransmissions, path MTU, SACK (SACK is enabled by default for iSCSI TCP configurations), window management (The iSCSI defaults are max-bandwidth = 1G, min-available-bandwidth = 70 Mbps, and round-trip-time = 1 ms.), buffer size (default send buffer size for iSCSI is 4096 KB), window congestion (enabled by default and the default burst size is 50 KB.), and maximum delay jitter (enabled by default and the default time is 500 microseconds.).
  - QoS—QoS configurations differ for iSCSI and FCIP interfaces.
- Identification of dynamic iSCSI initiator—iSCSI initiators are identified based on their IQN name or their IP address. In the absence of any configuration for the initiator (WWN or VSAN membership), the identifier key is the default connection. By default, the key is the IQN name but can be changed to IP address by toggling this mode.
- Proxy or transparent Initiator—For each iSCSI initiator with iSCSI target sessions, the switch creates a virtual FC initiator with a distinct pair of WWNs per VSAN. For targets that have access control per LUN, the WWN pair of each FC initiator must be configured in the target. The proxy initiator mode can be enabled to facilitate this configuration, in this case, all iSCSI initiators that connect to this iSCSI interface inherit the same WWN pair and create only one virtual FC initiator in each VSAN.

## Setting the QoS Values

To set the QoS values, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | In Fabric Manager, select <b>Interfaces &gt; FC Logical</b> from the Physical Attributes pane. You see the Interface tables in the Information pane.<br>In Device Manager, select <b>Interfaces &gt; Ethernet or iSCSI</b> . You see the interfaces dialog box. |
| <b>Step 2</b> | Click the <b>iSCSI TCP</b> tab. You see the iSCSI TCP configuration table.  |
| <b>Step 3</b> | Set the QoS field from 1 to 6.  |
| <b>Step 4</b> | Click the <b>Apply Changes</b> icon in Fabric Manager or click <b>Apply</b> in Device Manager to save these changes, or click <b>Undo Changes</b> in Fabric Manager or click <b>Cancel</b> in Device Manager to discard changes.                                |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## iSCSI Forwarding Mode

The iSCSI gateway on the IPS module has two modes of forwarding operation:

- The **pass-thru** mode (default): The IPS port converts an iSCSI PDU into an FCP frame or vice versa and then forwards it one frame or PDU at a time. The absence of buffering PDUs or frames keeps the operation latency low. To operate in this mode, the IPS port has to negotiate with its peers a suitable maximum size of the data payload in each frame/PDU. This is done during iSCSI login and FC PLOGI and the value is restricted by the TCP connection's maximum segment size (MSS) and the maximum Fibre Channel data payload size specified by the FC target. This usually results in a smaller maximum payload size than most hosts expect, thus comes the second mode of forwarding.
- The **store-and-forward** mode: The iSCSI client sends and receives an iSCSI data payload at the size it desires. This sometimes results in better performance for the client. The IPS port stores each TCP segment it receives until one full iSCSI PDU is received before converting and forwarding it as Fibre Channel frames to the FC target. In the opposite direction, the IPS port assembles all FC data frames of an exchange to build one iSCSI data-in PDU before forwarding it to the iSCSI client. The limitation on this mode is that the iSCSI CRC data digest cannot be used.

## iSCSI High Availability

The following high availability features are available for iSCSI configurations:

- [Multiple IPS Ports Connected to the Same IP Network, page 20-21](#)
- [VRRP-Based High Availability, page 20-22](#)
- [Ethernet PortChannel-Based High Availability, page 20-23](#)

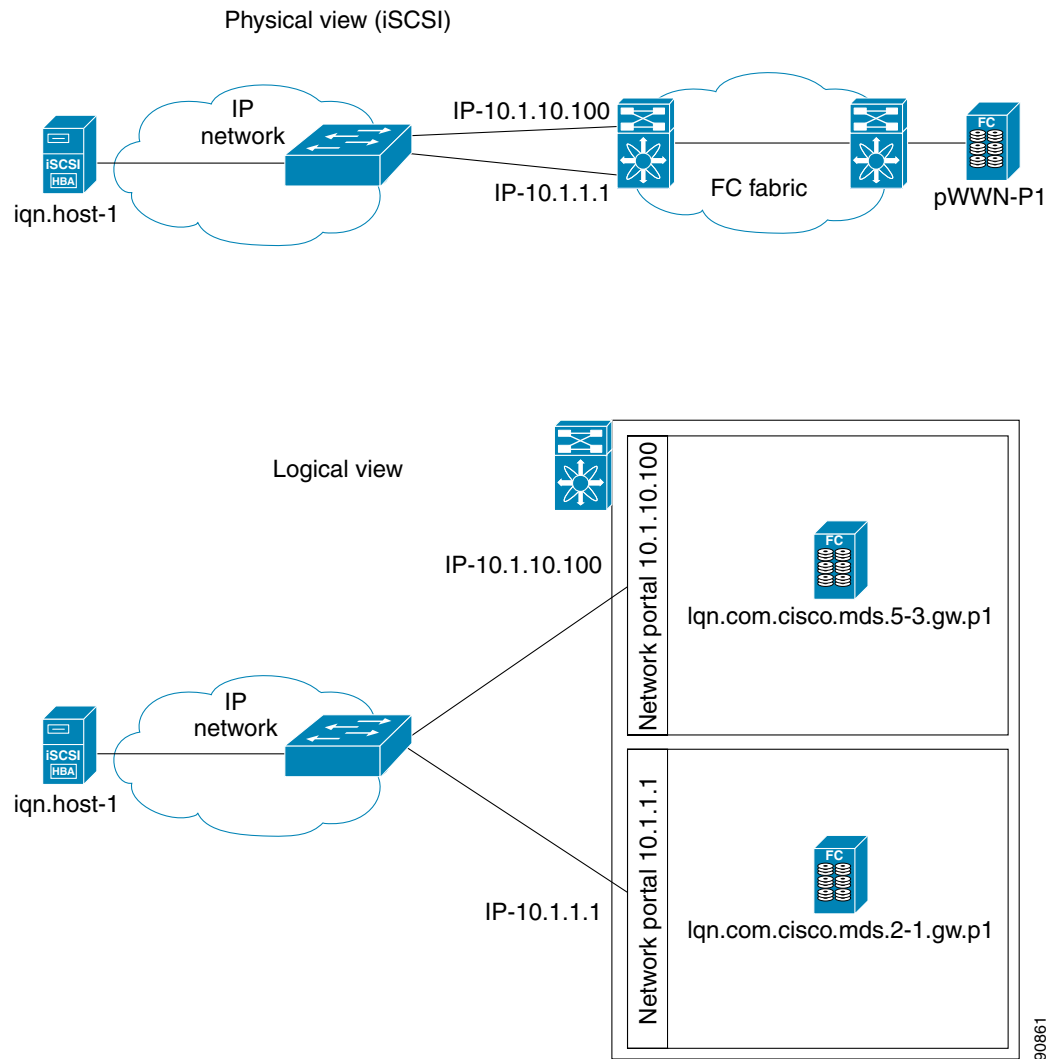


***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

### Multiple IPS Ports Connected to the Same IP Network

Figure 20-12 provides an example of a configuration with multiple Gigabit Ethernet interfaces in the same IP network.

**Figure 20-12 Multiple Gigabit Ethernet Interfaces in the Same IP Network**



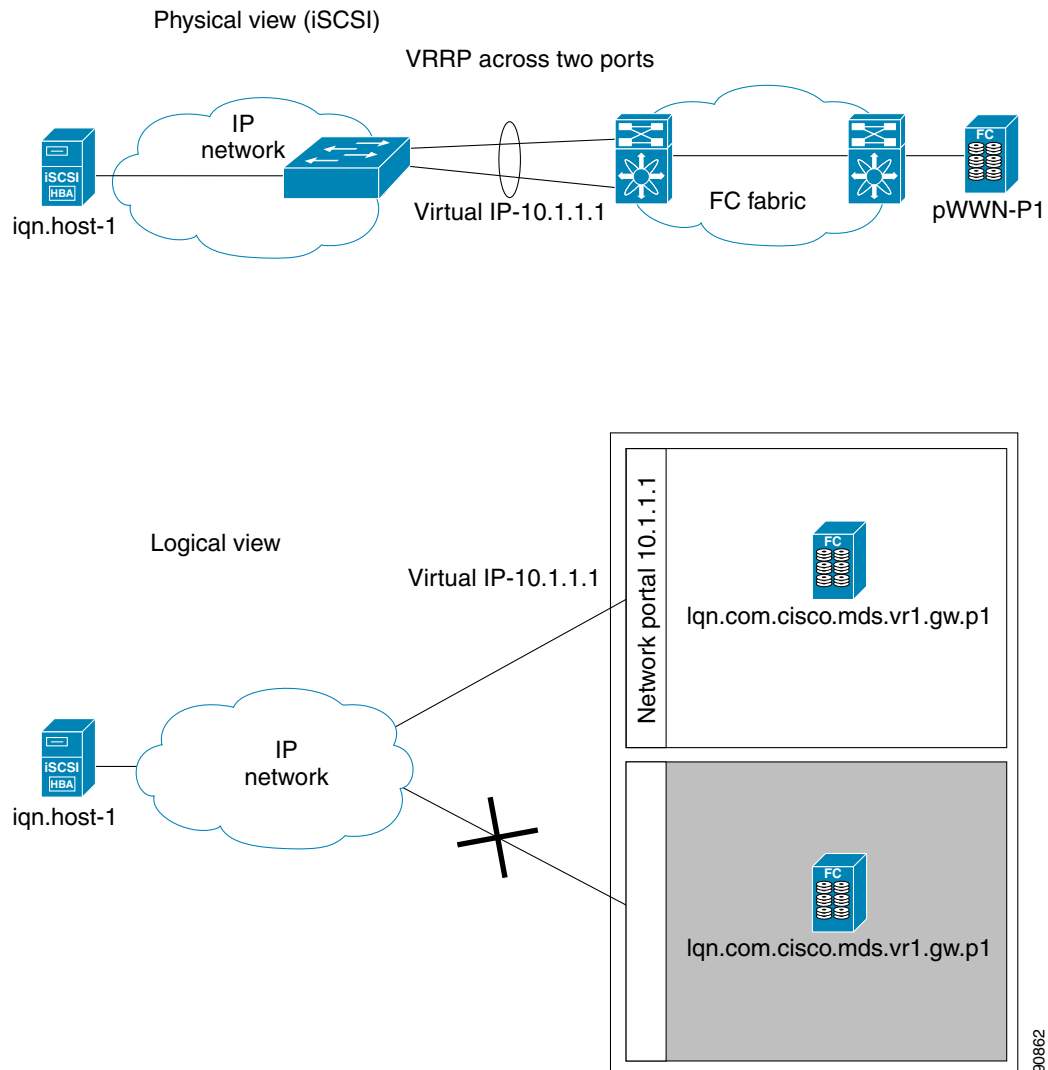
In Figure 20-12, each iSCSI host discovers two iSCSI targets for every physical Fibre Channel target (with different names). The multi-pathing software on the host provides load-balancing over both paths. If one Gigabit Ethernet interface fails, the host multi-pathing software is not affected because it can use the second path.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## VRRP-Based High Availability

Figure 20-13 provides an example of a VRRP-based high availability iSCSI configuration.

**Figure 20-13 VRRP-Based iSCSI High Availability**



In Figure 20-13, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. When the Gigabit Ethernet interface of the VRRP master fails, the iSCSI session is terminated. The host then reconnects to the target and the session comes up because the second Gigabit Ethernet interface has taken over the virtual IP address as the new master.



### Tip

Ports that act as VRRP master and backup can be on different switches. If you have a static WWN configuration for iSCSI initiators (see the [“Presenting iSCSI Hosts as Virtual Fibre Channel Hosts”](#) section on page 20-11), configure a different WWN for the iSCSI initiator for each switch. If you use a proxy initiator, be sure to configure a different pWWN on each iSCSI interface for each VRRP port used.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Ethernet PortChannel-Based High Availability

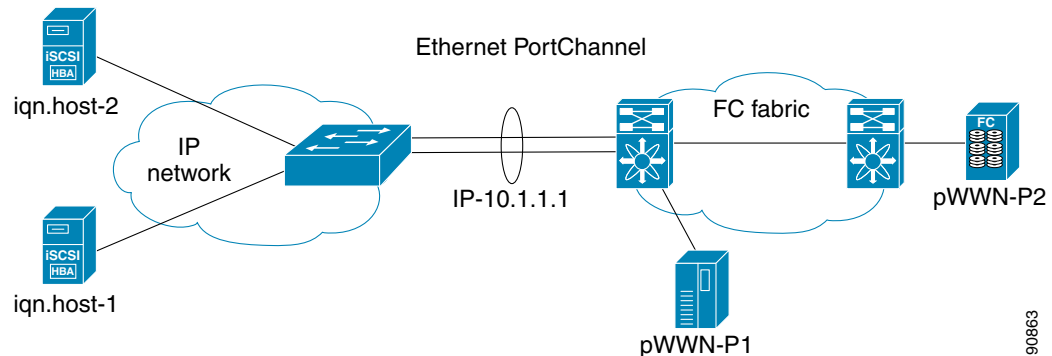


### Note

All iSCSI data traffic for one iSCSI link is carried on one TCP connection. Consequently, the aggregated bandwidth is 1 Gbps for that iSCSI link.

Figure 20-14 provides a sample Ethernet PortChannel-based high availability iSCSI configuration.

**Figure 20-14 Ethernet PortChannel-Based iSCSI High Availability**



In Figure 20-14, each iSCSI host discovers one iSCSI target for every physical Fibre Channel target. The iSCSI session from the iSCSI host to the virtual iSCSI target (on the IPS port) uses one of the two physical interfaces (because an iSCSI session uses one TCP connection). When the Gigabit Ethernet interface fails, the IPS module and the Ethernet switch transparently forwards all the frames on to the second Gigabit Ethernet interface.

## Configuring Storage Name Services

As of Cisco MDS SAN-OS Release 1.3(1), the Internet Storage Name Service (iSNS) client feature is available in all switches in the Cisco MDS 9000 Family with IPS modules installed.

iSNS services allow your existing TCP/IP networks to function more effectively as storage area networks by automating the discover and management of iSCSI devices. To facilitate these functions, the iSNS client functionality registers iSCSI portals and all targets accessible through a particular interface with an external iSNS server.



### Note

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is clicked, the other Information pane tabs that use CFS are activated.

## Creating iSNS Profiles and Tagging Profiles

The iSNS client functionality on each interface (Gigabit Ethernet interface or subinterface or PortChannel) registers information with its configured iSNS server using an iSNS profile. This process is referred to as tagging an iSNS profile to an interface. Each iSNS profile keeps information about an iSNS server IP address. One profile can be tagged to one or more interfaces.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Once a profile is tagged to an interface, the MDS switch opens a TCP connection to the iSNS server IP address (using a well-known iSNS port number 3205) in the profile and registers network entity and portal objects. It goes through the FC name server database and configuration to find storage nodes to register with the server.

Statically mapped virtual targets are registered if the associated Fibre channel pWWN is present in the FC name server database and no access control configuration prevents it. A dynamically mapped target is registered if the dynamic target importing is enabled. See the [“Using the iSCSI Wizard” section on page 20-5](#).

A storage node is deregistered from the iSNS server when it becomes unavailable either because of configuration changes (such as access control change or dynamic import disabling) or when the Fibre Channel storage port goes off-line. It is registered again when the node is online.

When the iSNS client is unable to register/deregister objects with the iSNS server (for example, the client is unable to make a TCP connection to the iSNS server), it retries every minute to re-register all iSNS objects for the affected interface(s) with the iSNS server.

Untagging a profile causes the network entity and portal to deregister from that interface.

## Creating an iSNS Profile

To create an iSNS profile, follow these steps:

- 
- Step 1** In Fabric Manager, select **End Devices > iSCSI** from the Physical Attributes pane. You see the iSCSI tables in the Information pane.  
In Device Manager, select **IP > iSCSI**. You see the iSCSI dialog box.
  - Step 2** Click the **iSNS Profiles** tab. You see the iSCSI authentication configuration table.
  - Step 3** Set the Name and IP address fields of the iSNS server.
  - Step 4** Click the **Apply Changes** icon in Fabric Manager or click **Create** in Device Manager to save these changes, or click **Undo Changes** in Fabric Manager or click **Cancel** in Device Manager to discard changes.
- 

## Modifying an iSNS Profile

To modify (tag) the iSNS profile for an interface, untag the interface from the currently tagged iSNS profile and then tag to a new iSNS profile.

To tag an interface to a profile using Device Manager, follow these steps:

- 
- Step 1** Select **Interfaces > Ethernet or iSCSI**. You see the interfaces dialog box.
  - Step 2** Click the **General** tab. You see the General interface configuration table.
  - Step 3** Select the **iSNS ProfileName**.
  - Step 4** Click **Apply** to save these changes, or click **Cancel** to discard changes.
-



## FICON Configuration

---

Fibre Connection (FICON) interface capabilities enhance the Cisco MDS 9000 Family by supporting both open systems and mainframe storage network environments. Inclusion of Control Unit Port (CUP) support further enhances the MDS offering by allowing in-band management of the switch from FICON processors.

The fabric binding feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. The Registered Link Incident Report (RLIR) application provides a method for a switchport to send a LIR to a registered Nx-port.



### Note

FICON features can be implemented in any switch in the Cisco MDS 9000 Family running Cisco MDS SAN-OS Release 1.3(x) or earlier. While no hardware changes are required, you do need the MAINFRAME\_PKG license to configure FICON parameters (see [Chapter 9, “Obtaining and Installing Licenses”](#)).

This chapter includes the following sections:


- [About FICON, page 21-2](#)
- [Enabling FICON, page 21-10](#)
- [Configuring FICON Ports, page 21-14](#)
- [FICON Configuration Files, page 21-16](#)
- [Port Swapping, page 21-18](#)
- [Clearing FICON Device Allegiance, page 21-19](#)
- [CUP In-Band Management, page 21-20](#)
- [Fabric Binding Configuration, page 21-20](#)
- [Displaying RLIR Information, page 21-25](#)
- [Calculating FICON Flow Load Balance, page 21-25](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## About FICON

The Cisco MDS 9000 Family supports the Fibre Channel Protocol (FCP), FICON, iSCSI, and FCIP capabilities within a single, high availability platform. This solution simplifies purchasing, reduces deployment and management costs, and reduces the complex evolution to shared mainframe and open systems storage networks (see [Figure 21-1](#)).

**Figure 21-1 Shared System Storage Network**



FCP and FICON are different FC4 protocols and their traffic are independent of each other. If required, devices using these protocols can be isolated using VSANs.

## MDS-Specific FICON Advantages

This section explains the additional FICON advantages in Cisco MDS switches.

### Fabric-Optimization with VSANs

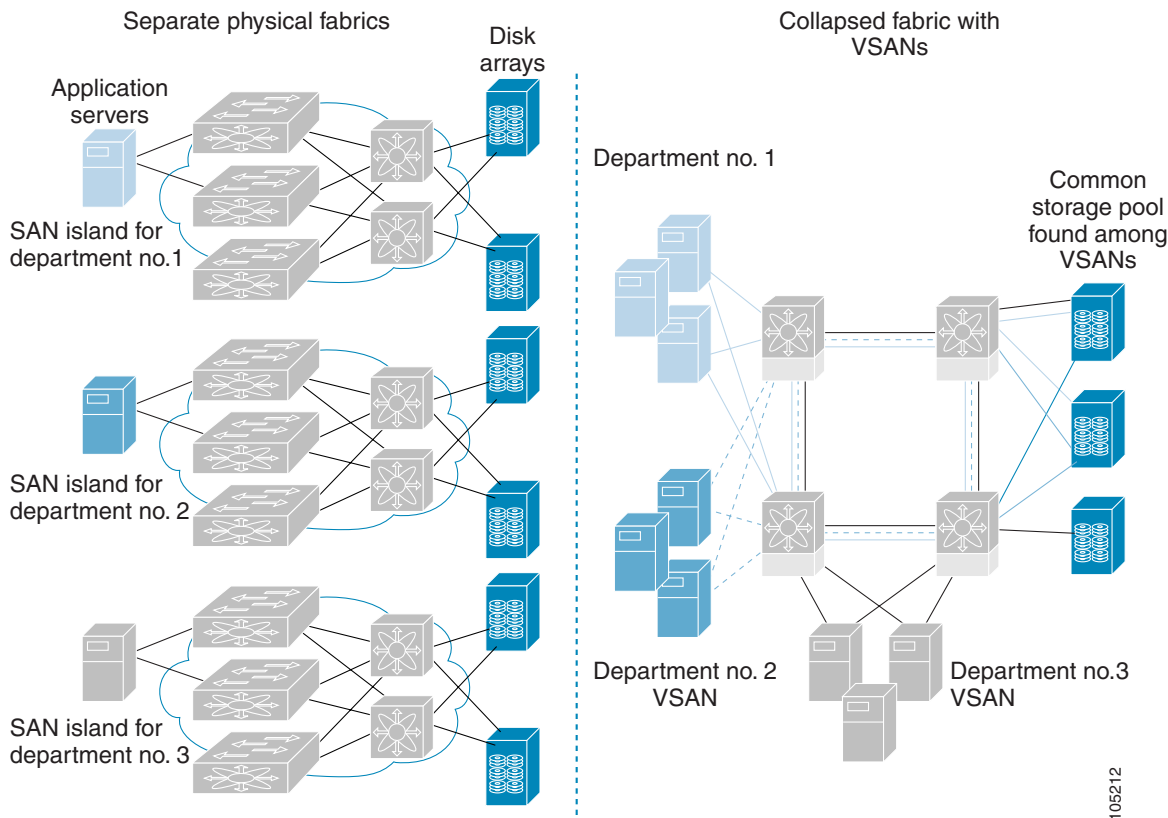
Generally, separate physical fabrics have a high level of switch management and have a higher implementation cost. Further, the ports in each island may be over-provisioned depending on the fabric configuration.

By using the Cisco MDS-specific VSAN technology, you can introduce greater efficiency between these physical fabrics by lowering the cost of over-provisioning and reducing the number of switches to be managed.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

VSANs also help you to move unused ports nondisruptively and provide a common redundant physical infrastructure (see [Figure 21-2](#)).

**Figure 21-2 VSAN-Specific Fabric Optimization**



VSANs enable global SAN consolidation by allowing you to convert existing SAN islands into virtual SAN islands on a single physical network. It provides hardware-enforced security and separation between applications or departments to allow coexistence on a single network. It also allows virtual rewiring to consolidate your storage infrastructure. You can move assets between departments or applications without the expense and disruption of physical relocation of equipment.



**Note**

While you can configure up to 256 VSANs in any Cisco MDS switch, you can enable FICON in eight of these VSANs.

## FCIP Support

The multilayer architecture of the Cisco MDS 9000 Family enables a consistent feature set over a protocol-agnostic switch fabric. Cisco MDS 9500 Series and 9200 Series switches transparently integrate Fibre Channel, FICON, and Fibre Channel over IP (FCIP) in one system. The FICON over FCIP feature enables cost-effective access to remotely located mainframe resources. With the Cisco

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

MDS 9000 Family platform, storage replication services such as IBM PPRC and XRC can be extended over metro to global distances using ubiquitous IP infrastructure and simplifying business continuance strategies.

**Caution**

When write-acceleration is enabled in an FCIP interface, a FICON VSAN will not be enabled in that interface. Likewise, if a FCIP interface is up in a FICON VSAN, write-acceleration cannot be enabled on that interface.

## PortChannel Support

The Cisco MDS implementation of FICON provides support for efficient utilization and increased availability of inter-switch links necessary to build stable large-scale SAN environments. PortChannels ensure an enhanced ISL availability and performance in Cisco MDS switches.

See [Chapter 17, “PortChannel Configuration,”](#) for more information on PortChannels.

## VSANs for FICON and FCP Intermixing

Cisco MDS 9000 Family FICON-enabled switches simplify deployment of even the most complex intermix environments. Multiple logical FICON, Z-Series Linux/FCP, and Open-Systems FCP fabrics can be overlaid onto a single physical fabric by simply creating VSANs as required for each service. VSANs provide both hardware isolation and protocol specific fabric services, eliminating the complexity and potential instability of zone-based intermix schemes.

By default, the FICON feature is disabled in all switches in the Cisco MDS 9000 Family. When the FICON feature is disabled, FC IDs can be allocated seamlessly. Intermixed environments are addressed by the Cisco MDS SAN-OS software. The challenge of mixing Fibre Channel Protocol (FCP) and FICON protocols are addressed by Cisco MDS switches when implementing VSANs.

Switches and directors in the Cisco MDS 9000 Family support FCP and FICON protocol intermixing at the port level. If these protocols are intermixed in the same switch, you can use VSANs to isolate FCP and FICON ports.

**Tip**

When creating an intermix environment, place all FICON devices in one VSAN (other than the default VSAN) and segregate the FCP switch ports in a separate VSAN (other than the default VSAN). This isolation ensures proper communication for all connected devices.

## Cisco MDS-Supported FICON Features

The Cisco MDS 9000 Family FICON features include:

- Flexibility and investment protection—The Cisco MDS 9000 Family shares common switching and service modules across the Cisco MDS 9500 Series and the 9200 Series.

Refer to the *Cisco MDS 9500 Series Hardware Installation Guide* and the *Cisco MDS 9200 Series Hardware Installation Guide*.

- High-availability FICON-enabled director—The Cisco MDS 9500 Series combines nondisruptive software upgrades, stateful process restart and failover, and full redundancy of all major components for a new standard in director-class availability. It supports up to 224 autosensing, 2/1-Gbps, FICON



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

or Fibre Channel FCP ports in any combination in a single chassis and up to 768 Fibre Channel ports in a single rack. The 1.44 Tbps of internal system bandwidth ensures smooth integration of future 10-Gbps modules.


- Infrastructure protection—Common software releases infrastructure protection is available across all Cisco MDS 9000 platforms.
- VSAN technology—The Cisco MDS 9000 Family introduces VSAN technology for hardware-enforced, isolated environments within a single physical fabric for secure sharing of physical infrastructure and enhanced FICON intermix support.
- Port-level configurations—BB\_credits, beacon mode, and port security for each port.
- Alias name configuration—instead of the WWN, for switches and attached node devices.
- Comprehensive security framework—The Cisco MDS 9000 Family supports RADIUS authentication, Simple Network Management Protocol Version 3 (SNMPv3), role-based access control, Secure Shell Protocol (SSH), Secure File Transfer Protocol (SFTP), VSANs, hardware-enforced zoning, ACLs, fabric binding, Fibre Channel Security Protocol (FC-SP), LUN zoning, read-only zones, and VSAN-based access control. See [and](#)
- View the local accounting log to locate FICON events.
- Unified storage management—Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console. See the [.](#)
- Port address-based configurations—port name, blocked or unblocked state, and the prohibit connectivity attributes. See the [.](#)
- Display the following information:
  - Individual Fibre Channel ports, such as the port name, port number, Fibre Channel address, operational state, type of port, and login data.
  - Nodes attached to ports.
  - Port performance and statistics.See the [section](#) in this chapter.
- Store and apply configuration files.
- FICON and Open Systems Management Server features if installed.
- Enhanced Cascading Support.
- Set the date and time on the switch.
- Configure SNMP trap recipients and community names.
- Call Home configurations—director name, location, description, and contact person.
- Configure preferred domain ID, FC ID persistence, and principle switch priority.
- Sophisticated SPAN diagnostics—The Cisco MDS 9000 Family provides industry-first intelligent diagnostics, protocol, decoding, and network analysis tools as well as integrated call-home capability for added reliability, faster problem resolution, and reduced service costs.
- Configure R\_A\_TOV, E\_D\_TOV.
- Perform maintenance tasks for the director including maintaining firmware levels, accessing the director logs, and collecting data to support failure analysis.
- Display and clear port-level incident alerts. [.](#)

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## FICON Port Numbering

With reference to the FICON feature, ports in Cisco MDS switches are identified by a statically defined 8-bit value known as the *port number*. Port numbers are assigned based on the module and the slot in the chassis. Port numbers cannot be changed and the first port in a switch always starts with a 0 (see [Figure 21-3](#)).

**Figure 21-3** Port Number in the Cisco MDS 9000 Family



The FICON port number is assigned based on the front panel location of the port and is specific to the slot in which the module resides. Even if the module is a 16-port module, 32-port numbers are assigned to that module—regardless of the module type (16-port or 32-port), the module’s physical presence in the chassis, or the port status (up or down).



### Note

Only Fibre Channel, PortChannel, and FCIP ports are mapped to FICON port numbers. Other types of interfaces do not have a corresponding port number.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Table 21-1 lists the port number assignment for the Cisco MDS 9000 Family of switches and directors.

**Table 21-1 FICON Port Numbering in the Cisco MDS 9000 Family**

Product	Slot Number	Implemented Port Allocation		Unimplemented Ports	Notes
		To Ports	To PortChannel/FCIP		
Cisco MDS 9200 Series	Slot 1	0 through 31	64 through 89	90 through 253 and port 255	Similar to a switching module.
	Slot 2	32 through 63			The first 16 port numbers in a 16-port module are used and the rest remain unused.
Cisco MDS 9506 Director	Slot 1	0 through 31	128 through 153	154 through 253 and port 255	
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			Supervisor modules are not allocated port numbers.
Cisco MDS 9509 Director	Slot 1	0 through 31	224 through 249	250 through 253 and port 255	The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 2	32 through 63			
	Slot 3	64 through 95			
	Slot 4	96 through 127			
	Slot 5	None			
	Slot 6	None			Supervisor modules are not allocated port numbers.
	Slot 7	128 through 159			The first 16 port numbers in a 16-port module are used and the rest remain unused.
	Slot 8	160 through 191			
	Slot 9	192 through 223			

## FICON Port Numbering Guidelines

The following guidelines apply to FICON port numbers:

- Supervisor modules do not have port number assignments.
- Port numbers are VSAN independent and do not change based on VSANs or TE ports.
- Each PortChannel must be explicitly associated with a FICON port number.
- When the port number for a physical PortChannel becomes uninstalled, the relevant PortChannel configuration is applied to the physical port.
- Each FCIP tunnel must be explicitly associated with a FICON port number. If the port numbers are not assigned for PortChannels or for FCIP tunnels, the associated ports will not come up.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## FCIP and PortChannel Port Numbers

FCIP and PortChannels cannot be used in a FICON-enabled VSAN unless they are explicitly bound to a port number.

## Port Addresses

By default, port numbers are the same as port addresses (see the [“Editing FICON Configuration Files” section on page 21-17](#)).

### Implemented and Unimplemented Port Addresses

An implemented port refers to any port address that is available in the chassis.

An unimplemented port refers to any port address that is not available in the chassis.



**Tip**

An unimplemented port is prohibited from communicating with an implemented port in a FICON setup and cannot be configured.

## Installed and Uninstalled Ports

An installed port refers to a port for which all required hardware is present. A specified port number in a VSAN can be implemented, and yet not installed, if any of the following conditions apply:

- The module is not present—for example, if module 1 is not physically present in slot 1 in a Cisco MDS 9509 Director, ports 0 to 31 are considered uninstalled.
- The small form-factor pluggable (SFP) port is not present—for example, if a 16-port module is inserted in slot 2 in a Cisco MDS 9509 Director, ports 48 to 63 are considered uninstalled.
- The port is not in a FICON-enabled VSAN—for example, if port 4 (of a 16-port module in slot 1) is configured in FICON-enabled VSAN 2, then only port 4 is installed and ports 0 to 3 and 5 to 15 are uninstalled—even if they are implemented in VSAN 2.

Another scenario is if VSANs 1 through 5 are FICON-enabled, and trunking-enabled interface fc1/1 has VSANs 3 through 10, then port address 0 is uninstalled in VSAN 1 and 2.

- The port is part of a PortChannel—for example, if interface fc 1/1 is part of PortChannel 5, port address 0 is uninstalled in all FICON VSANs.

## FC ID Allocation

FICON requires a predictable and static FC ID allocation scheme. When FICON is enabled, the FC ID allocated to a device is based on the port address of the port to which it is attached. The port address forms the middle byte of the fabric address. Additionally, the last byte of the fabric address should be the same for all devices in the fabric. By default, the last byte value is 0 and can be configured.




**Note**

You cannot configure persistent FC IDs in FICON-enabled VSANs.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Cisco MDS switches have a dynamic FC ID allocation scheme. When FICON is enabled or disabled on a VSAN, all the ports are flapped to switch from the dynamic to static FC IDs and vice versa (see [Figure 21-4](#)).

**Figure 21-4 Static FC ID Allocation for FICON**



Static FC ID allocation for interface fc3/5 includes the static domain ID (5), the port address (0x44), and the last byte value (0).

113134

## FICON Cascading

The Cisco MDS SAN-OS software allows multiple switches in a FICON network. To configure multiple switches, you must enable and configure fabric binding in that switch.

## FICON VSAN Prerequisites

To ensure that a FICON VSAN is operationally up, be sure to verify the following requirements:

- Set the default zone to permit, if you are not using the zoning feature. See the [“The Default Zone” section on page 15-14](#).
- Enable in-order delivery on the VSAN.
- Enable (and if required, configure) fabric binding on the VSAN.
- Verify that conflicting persistent FC IDs do not exist in the switch.
- Verify that the configured domain ID and requested domain ID match.
- Add the CUP (area FE) to the zone, if you are using zoning.

If any of these requirements are not met, the FICON feature cannot be enabled.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Enabling FICON

By default FICON is disabled in all switches in the Cisco MDS 9000 Family. When you enable the FICON feature in Cisco MDS switches, the following apply:

- You cannot disable in-order delivery for the FICON-enabled VSAN.
- You cannot disable fabric binding or static domain ID configurations for the FICON-enabled VSAN.
- The load balancing scheme is changed to Source ID (SID)—Destination ID (DID). You cannot change it back to SID—DID—OXID.
- The IPL configuration file is automatically created.

## Creating FICON VSANs and enabling FICON

When a new FICON VSAN is created, static (insistent) domain IDs, in-order delivery, and fabric binding must be enabled so the FICON VSAN can operate. When you enable the FICON feature in Cisco MDS switches, the following apply:

- The IPL configuration file is automatically created.
- You cannot disable in-order delivery, fabric binding, or static (insistent) domain ID configurations.

To create a FICON VSAN in Fabric Manager, follow these steps:

- 
- Step 1** In Fabric Manager, right-click **All VSANs** in the Logical pane, and click **Create VSAN**. You see the Create VSAN dialog box.
  - Step 2** Select the switches you want to be in the VSAN.
  - Step 3** Enter a VSAN ID.
  - Step 4** Enter the name of the VSAN, if desired.
  - Step 5** Select the type of load balancing, the interop value, and the administrative state for this VSAN.
  - Step 6** Check the **FICON** check box.
  - Step 7** To enable fabric binding for the selected switches, check that check box.
  - Step 8** Click **Create** to create the new VSAN, or click **Close** to close the dialog without creating the VSAN.
  - Step 9** Open Device Manager for each switch in the FICON VSAN.
  - Step 10** Select **VSANs** from the FC menu. You see the VSANs dialog box.
  - Step 11** Enter the VSAN membership information.
  - Step 12** Click the VSAN you want to become a FICON VSAN and select **Add** from the FICON drop-down list.
  - Step 13** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.
- 

To create a FICON VSAN in Device Manager, follow these steps:

- 
- Step 1** Choose **FC > VSANs**. You see the VSANs configuration dialog box.
  - Step 2** Click **Create VSAN**. You see the Create VSAN dialog box.
  - Step 3** Enter a VSAN ID.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Enter the name of the VSAN, if desired.
  - Step 5** Select the type of load balancing, the interop value, and the administrative state for this VSAN.
  - Step 6** Check the **FICON** check box.
  - Step 7** To enable fabric binding for the selected switches, check that check box.
  - Step 8** Click **Create** to create the FICON VSAN, or click **Close** to close the dialog without creating the FICON VSAN.
- 

## Deleting FICON VSANs

To delete a FICON VSAN in Fabric Manager, follow these steps:

- Step 1** Choose **All VSANS**. You see the VSAN table in the Information pane.
- Step 2** Click anywhere in the row for the VSAN which you want to delete.
- Step 3** Click the **Delete Row** icon to delete the VSAN.



**Note** Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.

---

To delete a FICON VSAN in Device Manager, follow these steps:

- Step 1** Choose **FICON > VSANs**. You see the VSAN dialog box.
- Step 2** Click the VSAN you want to disable FICON on.
- Step 3** Select **Remove** from the FICON drop-down list.
- Step 4** Click **Apply** to disable FICON on this VSAN or click **Close** to close the dialog box without making any changes.



**Note** Deleting the VSAN will also delete the associated FICON configuration file, and the file cannot be recovered.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Viewing FICON Director History

To view FICON director history, follow these steps:

- 
- Step 1** In Device Manager, select **VSANs** from the **FICON** menu. You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab if it is not already displayed.
  - Step 3** Click anywhere in the row for the VSAN for which you want to configure port information.
  - Step 4** Click the **Director History** button to display a history of FICON-related changes to this switch.
- 

## The code-page Option

FICON strings are coded in Extended Binary-Coded Decimal Interchange Code (EBCDIC) format. Refer to your mainframe documentation for details on the code page options.

Cisco MDS switches support **international-5**, **france**, **brazil**, **germany**, **italy**, **japan**, **spain-latinamerica**, **uk**, and **us-canada** (default) EBCDIC format options.



### Tip

This is an optional configuration. If you are not sure of the EBCDIC format to be used, we recommend retaining the **us-canada** (default) option.

---

To modify the code-page option using Device Manager, follow these steps:

- 
- Step 1** Select **VSANs** from the **FICON** menu. You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab if it is not already displayed.
  - Step 3** Choose the code-page option from the drop-down menu in the CodePage field for the FICON VSAN you want to configure.
  - Step 4** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.
- 

## FC ID Last Byte



### Caution

If the FICON feature is configured in cascaded mode, the Cisco MDS Switches use ISLs to connect to other switches.

---

FICON requires the last byte of the fabric address to be the same for all allocated FC IDs. By default, this value is set to 0. You can only change the FC ID last byte when the FICON switch is in the offline state.



*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## FICON Host Control

By default, the clock in each VSAN is the same as the switch hardware clock. Each VSAN in a Cisco MDS switch represents a virtual director. The clock and time present in each virtual director can be different. To maintain separate clocks for each VSAN, the Cisco MDS SAN-OS software maintains the difference of the VSAN-specific clock and the hardware-based director clock. When a host (mainframe) sets the time, the Cisco MDS SAN-OS software updates this difference between the clocks. When a host reads the clock, it computes the difference between the VSAN-clock and the current director hardware clock and presents a value to the mainframe.

To allow the host (mainframe) to control the Cisco MDS switch using Device Manager, follow these steps:

- 
- Step 1** Select **VSANs** from the FICON menu. You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab if it is not already displayed.
  - Step 3** Check the **Offline Sw** check box under Host can to allow the mainframe to move a switch to the offline state.
  - Step 4** Check the **Sync Time** check box under Host can to allow the mainframe to set the system time on the switch.
  - Step 5** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.
- 

## Host Changes FICON Port Parameters

By default, mainframe users are not allowed to configure FICON parameters on Cisco MDS switches—they can only query the switch.

To allow the host (mainframe) to configure FICON parameters on the Cisco MDS switch using Device Manager, follow these steps:

- 
- Step 1** Select **VSANs** from the FICON menu. You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab if it is not already displayed.
  - Step 3** Check the **By Host** check box under Port Control to allow the mainframe to control a switch.
  - Step 4** Check the **By SNMP** check box under Port Control can to allow SNMP users to configure FICON on the switch.
  - Step 5** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.
- 



### Note

If you disable SNMP use in the Cisco MDS switch, you cannot configure FICON parameters using the Fabric Manager.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## FICON Information Refresh Note

When viewing FICON information through the Device Manager dialog boxes, you must manually refresh the display by clicking the Refresh button in order to see the latest updates. This is true whether you configure FICON through the CLI or through the Device Manager.

There is no automatic refresh of FICON information. This information would be refreshed so often that it would affect performance.

## Configuring FICON Ports

You can perform FICON configurations on a per-port address basis in the Cisco MDS 9000 Family of switches.

Even if a port is uninstalled, the port address-based configuration is accepted by the Cisco MDS switch. This configuration is applied to the port when the port becomes installed.

## Port Blocking

If you block a port, the port is retained in the operationally down state. If you unblock a port, a port initialization is attempted. When a port is blocked, data and control traffic are not allowed on that port.

Physical Fibre Channel port blocks will continue to transmit an Off-Line State (OLS) primitive sequence on a blocked port.



### Caution

---

You cannot block or prohibit the CUP port (0XFE).

---

If a port is shut down, unblocking that port does not initialize the port.

## Port Prohibiting

To prevent implemented ports from talking to each other, you can configure prohibits between two or more ports. If you prohibit ports, the specified ports are prevented from communicating with each other.



### Note

---

Unimplemented ports are always prohibited.

---



### Tip

---

You cannot prohibit a PortChannel or FCIP interface.

---

Prohibit configurations are always symmetrically applied—if you prohibit Port 0 from talking to port 15, port 15 is automatically prohibited from talking to port 0.



### Note

---

If an interface is already configured in E or TE mode and you try to prohibit that port, your prohibit configuration is rejected. Similarly, if a port is not up and you prohibit that port, the port is not allowed to come up in E mode nor in TE mode.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring Port Blocking and Port Prohibiting

To configure port blocking or port prohibiting for FICON using Device Manager, follow these steps:

- 
- Step 1** .Select **VSANs** from the FICON menu. You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab if it is not already displayed.
  - Step 3** Click the **Port Configuration**. You see the FICON Port Configuration dialog box.
  - Step 4** Set the port block and prohibit configuration for the selected FICON VSANs.
  - Step 5** Click **Apply** to save these changes or click **Close** to exit the dialog box without saving changes.
- 

## Entering FICON Port Configuration Information



### Note

To view the latest FICON information, you must click the **Refresh** button. See the “[FICON Information Refresh Note](#)” section on page 21-14.

To display FICON port configuration information, follow these steps:

- 
- Step 1** In Device Manager, select **VSANs** from the FICON menu.  
You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab.
  - Step 3** Click anywhere in the row for the VSAN for which you want to configure port information.
  - Step 4** Click **Port Configuration** to display the Port Configuration dialog box.
  - Step 5** Enter the Port Configuration information. Click **Apply** to save the configuration information, or click **Cancel** to exit the dialog without saving.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Viewing FICON Port Attributes



### Note

To view the latest FICON information, you must click the **Refresh** button. See the “[FICON Information Refresh Note](#)” section on page 21-14.

To view FICON port attributes, follow these steps:

- 
- Step 1** In Device Manager, select **VSANs** from the FICON menu.  
You see the FICON VSAN configuration dialog box.
  - Step 2** Click the **VSANs** tab.
  - Step 3** Click anywhere in the row for the VSAN for which you want to configure port information.
  - Step 4** Click **Port Attributes** to display the Port Attributes dialog box.
- 

## FICON Configuration Files

You can save up to 16 FICON configuration files on each FICON-enabled VSAN (in persistent storage). The file format is proprietary to IBM. These files can be read and written by IBM hosts using the in-band CUP protocol. Additionally, you can use the Cisco MDS CLI or Fabric Manager applications to operate these FICON configuration files.



### Note

Multiple FICON configuration files with the same name can exist in the same switch, provided they reside in different VSANs. For example, you can create a configuration file named XYZ in both VSAN 1 and VSAN 3.

When you enable the FICON feature in a VSAN, the switches always use the startup FICON configuration file, called IPL. This file is created with a default configuration as soon as FICON is enabled in a VSAN.



### Caution

When FICON is disabled on a VSAN, all the FICON configuration files are irretrievably lost.

FICON configuration files contain the following configuration for each implemented port address:

- Block
- Prohibit mask
- Port address name



### Note

Normal configuration files used by Cisco MDS switches include FICON-enabled attributes for a VSAN, port number mapping for PortChannels and FCIP interfaces, port number to port address mapping, port and trunk allowed VSAN configuration for ports, in-order guarantee, configuring static domain ID, and fabric binding configuration.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Accessing FICON Configuration Files

Only one user can access the configuration file at any given time:

- If this file is being accessed by user 1, user 2 cannot access this file.
- If user 2 does attempt to access this file, an error is issued to user 2.
- If user 1 is inactive for more than 15 seconds, the file is automatically closed and available for use by any other permitted user.

FICON configuration files can be accessed by any host, SNMP, or CLI user who is permitted to access the switch. The locking mechanism in the Cisco MDS SAN-OS software restricts access to one user at a time per file. This lock applies to newly created files and previously saved files. Before accessing any file, you must lock the file and obtain the file key. A new file key is used by the locking mechanism for each lock request. The key is discarded when the lock timeout of 15 seconds expires. The lock timeout value cannot be changed.

If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.



### Note

To view the latest FICON information, you must click the **Refresh** button. See the “[FICON Information Refresh Note](#)” section on page 21-14.

## Copying FICON Configuration Files

The Cisco MDS SAN-OS software maintains different configuration files to support a FICON network. These configuration files can be saved using the copy running-config startup-config command, or using Device Manager. FICON configuration files do not contain the following information that is normally saved with the running configuration:



### Note

To view the latest FICON information, you must click the **Refresh** button. See the [FICON Information Refresh Note, page 21-14](#) for more information.

- Port number to port address mapping
- PortChannel to port number mapping
- Port swap occurrences
- FICON enabled VSANs

FICON configuration files are independent of these parameters. Instead, this information is stored in persistent storage as they can be modified independent of the startup configuration.

## Editing FICON Configuration Files

The configuration file submode allows you to create and edit FICON configuration files. If a specified file does not exist, it is created. Up to 16 files can be saved. Each file name is restricted to eight alphanumeric characters.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Managing FICON Configuration Files In Device Manager

To manage a FICON file using Device Manager, follow these steps:

- 
- Step 1** Select **VSANS** from the **FICON** menu. You see the FICON VSANs dialog box.
  - Step 2** Click the **Files** tab.
  - Step 3** Click **Create** to create a new FICON configuration file.
    - a.** Enter the VSAN ID for the FICON VSAN you want to configure.
    - b.** Enter the file name and the description.
    - c.** Click **Create** to create the new file, or click **Close** to close the dialog without creating the file.
  - Step 4** Click **Copy** to copy the file to a new file.
  - Step 5** Click **Open** to edit the FICON configuration file.
  - Step 6** Click **Delete** to delete the FICON configuration file.
  - Step 7** Click **Apply** to apply the FICON configuration file.
- 

## Port Swapping

The FICON port swap feature is only provided for maintenance purposes.

The FICON port swapping feature causes all configuration associated with *old-port-number* and *new port-number* to be swapped, including VSAN configurations.

Cisco MDS switches allow port swapping for non-existent ports as follows:

- Only FICON-specific configurations (prohibit, block, and port address mapping) are swapped.
- No other system configuration is swapped.
- All other system configurations are only maintained for existing ports.



### Tip

If you check the **Active = Saved** check box on any FICON VSAN, then the swapped configuration is automatically saved to startup. Otherwise, you must explicitly save the running configuration immediately after swapping the ports.

---

Once you swap ports, the switch automatically performs the following actions:

- Shuts down both the old and new ports.
- Swaps the port configuration.
- If you attempt to bring the port up, you must explicitly shut down the port to resume traffic.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Port Swapping Guidelines

Be sure to follow these guidelines when using the FICON port swap feature:

- Port swapping is not supported for logical ports (PortChannels, FCIP links). Neither the *old-port-number* nor the *new-port-number* can be a logical port.
- Port swapping is not supported between physical ports that are part of a PortChannel. Neither the *old-port-number* nor the *new-port-number* can be a physical port that is part of a PortChannel.
- Before performing a port swap, the Cisco MDS SAN-OS software performs a compatibility check. If the two ports have incompatible configurations, the port swap is rejected with an appropriate reason code. For example, if a port with BB\_credits as 25 is being swapped with an OSM port for which a maximum of 12 BB\_credits is allowed (not a configurable parameter), the port swapping operation is rejected.
- If ports have default values (for some incompatible parameters), then a port swap operation is allowed and the ports retain their default values. If you swap a 16-port module with a 32-port module, the BB\_credits will no longer be compatible and the ports can be swapped. If BB\_credits are not configured, the default settings will still be in effect at the time of the swap.



### Note

The 32-port module guidelines also apply for port swapping configurations (see the “[32-Port Configuration Guidelines](#)” section on page 18-5).

## Swapping FICON Ports



### Note

To view the latest FICON information, you must click the **Refresh** button. See the “[FICON Information Refresh Note](#)” section on page 21-14 for more information.

To swap ports using Device Manager, follow these steps:

- Step 1** Select two Fibre Channel ports, by holding down the **CTRL** key and clicking on them with the mouse.
- Step 2** Select **Swap Selected Ports** from the FICON menu.

## Clearing FICON Device Allegiance

FICON requires serialization of access among multiple mainframes, CLI, and SNMP sessions be maintained on Cisco MDS 9000 Family switches by controlling device allegiance for the currently executing session. Any other session is denied permission to perform configuration changes unless the required allegiance is available.



### Caution

This task terminates the currently executing session.

Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).

# CUP In-Band Management

The Control Unit Port (CUP) protocol configures access control and provides unified storage management capabilities from a mainframe computer. Cisco MDS 9000 FICON-enabled switches are fully IBM CUP standard compliant for in-band management using the IBM S/A OS/390 I/O operations console.



Note

The CUP specification is proprietary to IBM.

CUP is supported by switches and directors in the Cisco MDS 9000 Family. The CUP function allows the mainframe to manage the Cisco MDS switches.

Host communication includes control functions such as blocking and unblocking ports, as well as monitoring and error reporting functions.

## Fabric Binding Configuration

The Cisco MDS SAN-OS Release 1.3(x) fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis and can only be implemented in FICON VSANs. You can still perform fabric binding configuration in a non-FICON VSAN—these configurations will only come into effect after FICON is enabled.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol in FICON networks to ensure that the list of authorized switches is identical in all switches in the fabric.

## Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other (see [Table 21-2](#)).

Table 21-2 Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWN and a persistent Domain ID.	Uses pWWNs/nWWNs or fWWNs/switch WWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN port(s). The switchport, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (list).
Activation is required on a per VSAN basis.	Activation is required on a per VSAN basis.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 21-2**      ***Fabric Binding and Port Security Comparison (continued)***

<b>Fabric Binding</b>	<b>Port Security</b>
User defines specific switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	User specifies the specific physical port(s) to which another device can connect.
Does not learn logging in switches.	Learns about switches or devices if in learning mode.

Port-level checking for xE-ports

- switch login uses both port binding as well as the fabric binding feature for a given VSAN.
- Binding checks are done on the port VSAN:
  - E-port security binding check is done on port VSAN.
  - TE-port security binding check is done in each allowed VSAN.

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

## Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. However, enforcement of fabric binding at the time of activation happens only if the VSAN is a FICON VSAN. The fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database.

To configure fabric binding in each switch in the fabric, follow these steps.

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Enable the fabric configuration feature.  |
| <b>Step 2</b> | Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric. |
| <b>Step 3</b> | Activate the fabric binding database.   |
| <b>Step 4</b> | Save the fabric binding configuration.  |
| <b>Step 5</b> | Verify the fabric binding configuration.  |
- 

## Enabling Fabric Binding

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family. The configuration and verification commands for the fabric binding feature are only available when fabric binding is enabled on a switch. When you disable this configuration, all related configurations are automatically discarded.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To enable fabric binding using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Fabric > VSANxxx > Fabric Binding** in the Logical Domain pane and then click the **Controls** tab in the Information pane.
  - Step 2** Set the Command drop-down menu to **enable** for the VSAN(s) on which you want to enable fabric binding.
  - Step 3** Click the **Apply Changes** icon in the Information pane to enable fabric binding.
- 

## Configuring a List of Switch WWNs In a Fabric

A user-specified fabric binding list contains a list of switch WWNs (sWWNs) within a fabric. If a sWWN attempts to join the fabric, and that sWWN is not in the list or the sWWN is using a domain ID that differs from the one specified in the allowed list, the ISL between the switch and the fabric is automatically isolated in that VSAN and the switch is denied entry into the fabric.

The persistent domain ID must be specified along with the sWWN. Domain ID authorization is required in FICON VSANs where the domains are statically configured and the end devices reject a domain ID change in all switches in the fabric.

To configure a list of switches for fabric binding using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Fabric > VSANxxx > Fabric Binding** in the Logical Domain pane and then click the **Config Database** tab in the Information pane.
  - Step 2** Click the **Create Row** icon to add a switch to the list of allowed switches for fabric binding.
  - Step 3** Click the **Apply Changes** icon in the Information pane to enable the fabric binding.
- 

## Activating Fabric Binding

The fabric binding maintains a configuration database (config-database) and an active database. The config-database is a read-write database that collects the configurations you perform. These configurations are only enforced upon activation. This activation overwrites the active database with the contents of the config database. The active database is read-only and is the database that checks each switch that attempts to log in.

By default, the fabric binding feature is not activated. You cannot activate the switch if entries existing in the config database conflict with the current state of the fabric. For example, one of the already logged in switches may be denied login by the config database. You can choose to forcefully override these situations.



### Note

After activation, any already logged in switch that violates the current active database will be logged out, and all switches that were previously denied login because of fabric binding restrictions are reinitialized.

The fabric binding feature must be enabled in each switch in the fabric that participates in the fabric binding. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the `forceActivate` option.

To activate fabric binding, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** in the Logical Domain pane and then click the **Actions** tab in the Information pane.
  - Step 2** Set the Action drop-down menu to **activate** or **forceActivate** for the VSAN(s) for which you want to activate fabric binding.
  - Step 3** Click the **Apply Changes** icon in the Information pane to activate the fabric binding.
- 

## Saving Fabric Binding Configurations

When you save the fabric binding configuration, the config database and the active database are both saved to the startup configuration and are available after a reboot.



### Caution

You cannot deactivate or disable fabric binding in a FICON-enabled VSAN.

## Deactivating Fabric Binding

To deactivate fabric binding, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** and then click the **Actions** tab in the Information pane.
  - Step 2** Set the Action drop-down menu to **deactivate** for the VSAN(s) for which you want to deactivate fabric binding.
  - Step 3** Click the **Apply Changes** icon to deactivate the fabric binding.
- 

## Fabric Binding CopyActive to Config

To copy the active fabric binding to the configuration file, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** in the Logical Domains pane and then click the **Actions** tab in the Information pane.
  - Step 2** Click the **CopyActive ToConfig** check box for the VSAN(s) for which you want to copy fabric binding.
  - Step 3** Click the **Apply Changes** icon to copy the fabric binding.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Creating a Fabric Binding Configuration

To create a fabric binding configuration, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** in the Logical Domains pane and then click the **Config Database** tab in the Information pane.
  - Step 2** Click **Create** to display the Config Database - Create dialog box.
  - Step 3** Enter the VSAN ID, the peer WWN, and the domain ID.
  - Step 4** Click the **Create Row** icon to create the fabric binding configuration.
- 

## Deleting a Fabric Binding Configuration

To delete a fabric binding configuration, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** in the Logical Domains pane and then click the **Config Database** tab in the Information pane.
  - Step 2** Click in the row for the VSAN for which you want to delete the fabric binding configuration.
  - Step 3** Click the **Delete Row** icon to delete the fabric binding configuration.
- 

## Viewing Fabric Binding Active Database

To view the fabric binding active database, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** and click the **Active Database** tab. You see the active database.
- 

## Viewing Fabric Binding Violations

To view fabric binding violations, follow these steps:

- 
- Step 1** In Fabric Manager, select **Fabric > VSANxxx > Fabric Binding** and click the **Violations** tab. You see the violations.
-

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Clearing Fabric Binding Statistics

To clear fabric binding statistics, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | In Fabric Manager, select <b>Fabric &gt; VSANxxx &gt; Fabric Binding</b> and click the <b>Statistics</b> tab.<br>You see the statistics in the Information pane. |
| <b>Step 2</b> | Check the <b>Clear</b> check box for the VSAN(s) for which you want to clear statistics.   |
| <b>Step 3</b> | Click the <b>Apply Changes</b> icon.   |
- 

## Viewing EFMD Statistics

To view EFMD statistics, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | In Fabric Manager, select <b>Fabric &gt; VSANxxx &gt; Fabric Binding</b> and click the <b>EFMD Statistics</b> tab. |
| <b>Step 2</b> | You see the EFMD statistics.   |
- 

## Displaying RLIR Information

The Registered Link Incident Report (RLIR) application provides a method for a switchport to send an LIR to a registered Nx-port. It is a highly-available application.

When a Link Incident Record (LIR) is detected in FICON-enabled switches in the Cisco MDS 9000 Family from a RLIR Extended Link Service (ELS), it sends that record to the members in its Established Registration List (ERL).

In case of multi-switch topology, a Distribute Registered Link Incident Record (DRLIR) Inter Link Service (ILS) is sent to all reachable remote domains along with the RLIR ELS. On receiving the DRLIR ILS, the switch extracts the RLIR ELS and sends to the members of the ERL.

The Nx-ports interested in receiving the RLIR ELS send Link Incident Record Registration (LIRR) ELS request to the management server on the switch. The RLIRs are processed on a per-VSAN basis.

To view RLIR information using Device Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>FICON &gt; RLIR ERL....</b> You see the Show RLIR ERL dialog box. |
| <b>Step 2</b> | Click <b>Close</b> to close the dialog box.                                 |
- 

## Calculating FICON Flow Load Balance

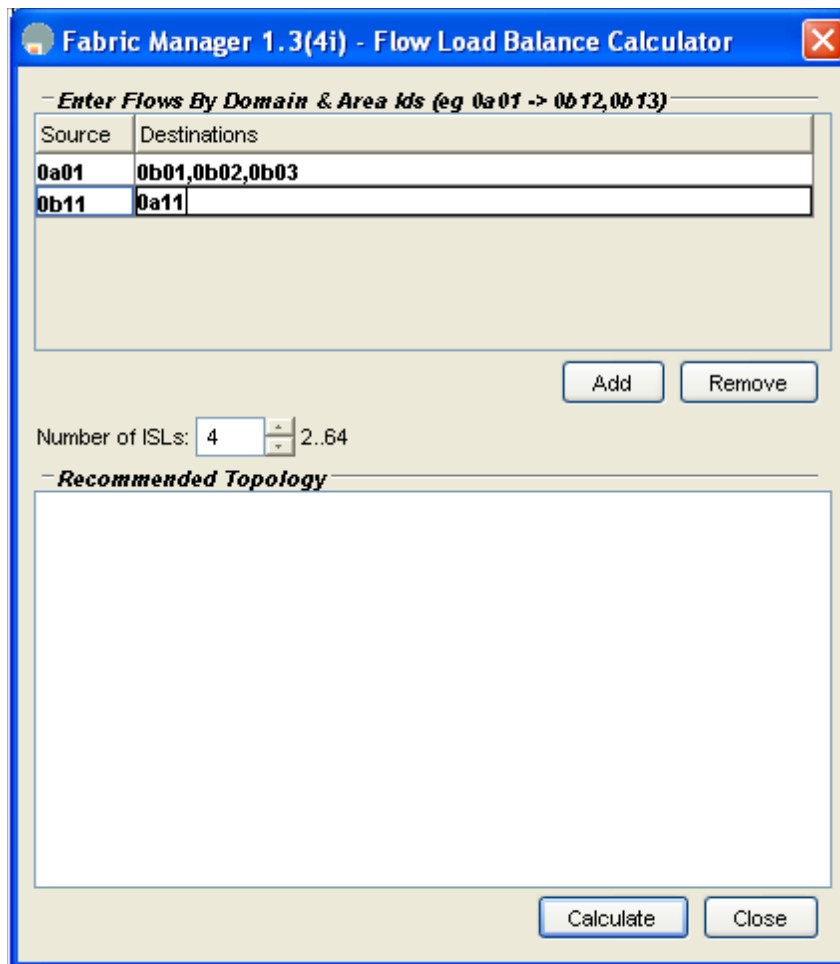
The FICON Flow Load Balance Calculator allows you to get the best load balancing configuration for your FICON flows. The calculator does not rely on any switch or flow discovery in the fabric. It is available from the Fabric Manager Tools menu.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

To use the FICON Flow Load Balance calculator, follow these steps:

- Step 1** Click **Tools > Other > FICON Flow Load Balance Calculator**. You see the Flow Load Balance Calculator (see [Figure 21-5](#)).
- Step 2** Click **Add** to enter the source and destination(s) flows.
- Use 2 byte hex (Domain and Area IDs) as shown in [Figure 21-5](#). You can copy and paste these IDs, and then edit them if you need to. To remove a row, select it and click **Remove**.

**Figure 21-5** *Flow Load Balance Calculator - Initial Screen*



- Step 3** Enter (or select) the number of ISLs between the two switches (for example, between Domain ID 0a and 0b in [Figure 21-5](#).)
- Step 4** Click **Calculate** to show the recommended topology

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

In the example shown in Figure 21-6, there are 12 ISLs between domains 0a and 0b with 13 flows. In this case, the best balance is 2 Port Channels with 2 members each, and 8 regular ISLs.

**Figure 21-6** Flow Load Balance Calculator - Example

**Fabric Manager 1.3(4i) - Flow Load Balance Calculator**

— Enter Flows By Domain & Area ids (eg 0a01 -> 0b12,0b13)

Source	Destinations
0a01	0b01,0b02,0b03
0b11	0a11,0af0
0a02	0b01,0b02,0b03
0b12	0a11,0a12
0a03	0b04,0b05,0b06

Add Remove

Number of ISLs: 12 2.64

— Recommended Topology

```
# 1: Port Channel with 2 ISL members
# 2: Port Channel with 2 ISL members
# 3: ISL
# 4: ISL
# 5: ISL
# 6: ISL
# 7: ISL
# 8: ISL
# 9: ISL
# 10: ISL
```

Calculate Close



**Note**

If you change flows or ISLs, you must click **Calculate** to see the new recommendation.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Configuring Intelligent Storage Services

---

The Storage Services Module (SSM) supports Intelligent Storage Services on Cisco MDS 9000 Family switches running Cisco MDS SAN-OS Release 2.0(2b) or later software. Intelligent Storage Services support in Cisco MDS SAN-OS Release 2.0(2b) includes the following:

- Fibre Channel write acceleration
- SCSI flow statistics

Intelligent Storage Services support in Cisco MDS SAN-OS Release 2.1(1) include the following:

- SANTap
- Network-Accelerated Serverless Backup (NASB)
- Third-party application support

This chapter includes the following sections:

- [Intelligent Storage Services, page 22-1](#)
- [SCSI Flow Services, page 22-3](#)
- [Fibre Channel Write Acceleration, page 22-4](#)
- [SCSI Flow Statistics, page 22-5](#)
- [SANTap, page 22-7](#)
- [NASB, page 22-11](#)

## Intelligent Storage Services

All Intelligent Storage Services must be enabled on an SSM before the service can be configured. For switches running Cisco MDS SAN-OS Release 2.1(1a) or later software, these services are enabled for all ports on the SSM, or provisioned in groups of four ports. Switches running earlier releases that support intelligent storage services enable a service across all ports.



### Note

---

The four port groups are contiguous, requiring you to configure ports 1 through 4, 5 through 8, and so on.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Intelligent Storage Services include the following:

- SCSI flow services
- SANTap
- NASB
- Third-party applications

## Enabling Intelligent Storage Services

In Cisco MDS SAN-OS Release 2.1(1a) or later, you can provision a subset of the ports for an SSM feature. The port range must be a multiple of four (for example fc4/1 through fc4-12).

To enable Intelligent Storage Services in Fabric Manager for an SSM and provision all ports or a group of ports to use these services, follow these steps:

- 
- Step 1** Choose **End Devices > SSM Features** in the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane.
- Step 2** Click the **SSM** tab. You see the set of configured services in the Information pane.
- Step 3** Click the **Create Row...** icon to enable a new service on an SSM. You see the SSM Create dialog box.
- Step 4** Select the **Switch** and **SSM Card** you want to configure.
- Step 5** Uncheck the **Use All Ports on Card** check box if you want to provision a subset of the ports on the card to use this service.
- Step 6** Select the port range you want to provision for using this service (starting port and ending port).




---

**Note** The port range must be a multiple of four (for example fc4/1 through fc4-12).

---

- Step 7** Select the Feature you want to enable on these ports from the drop-down list of services.
- Step 8** Set the **PartnerImageURI** field if you are enabling a third-party application that requires an image loaded onto the SSM.
- Step 9** Click **Create** to create this row and enable this service or click **Cancel** to discard all changes.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Disabling Intelligent Storage Services

To disable Intelligent Storage Services in Fabric Manager for an SSM and free up a group of ports that used these services, follow these steps:

- Step 1** Choose **End Devices > SSM Features** in the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane.
- Step 2** Click the **SSM** tab. You see the set of configured services in the Information pane.
- Step 3** Select the row in the table that you want to disable.
- Step 4** Check the **Reboot Card on Delete** check box if you want to force the card to reboot after disabling the service. This is equivalent to the CLI **force** option.
- Step 5** Click the **Delete Row** icon to delete this row. The ports that were provisioned for this service become available for provisioning in another service.



**Note** If **Reboot Card on Delete** was checked, then the SSM module reboots.

## SCSI Flow Services

A SCSI flow is a SCSI initiator/target combination. SCSI Flow Services provides enhanced features for SCSI flows such as write acceleration and flow monitor for statistics gathering on an SSM. SCSI flows can exist between an initiator/target combination on one SSM, or between two SSMs on different Cisco MDS switches.



**Note** For statistics monitoring, the target device is not required to be connected to an SSM.

The SCSI flow manager resides on a supervisor module and handles the configuration of SCSI flows, validating them and relaying configuration information to the appropriate SSM. It also handles any dynamic changes to the status of the SCSI flow due to external events. The SCSI flow manager registers events resulting from operations, such as port up or down, VSAN suspension, and zoning that affects the SCSI flow status, and updates the flow status and configuration accordingly.

The SCSI flow manager on the initiator communicates to its peer on the target side using Cisco Fabric Services (CFS). Peer communication allows the initiator SCSI flow manager to validate target parameters and program information on the target side.

Fabric Manager creates SCSI flows when Fibre Channel write acceleration or SCSI flow statistics gathering is configured.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring SCSI Flow Services

A SCSI flow specification consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN



### Note

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is click, the other tabs in the Information pane that use CFS are activated.

To configure a Fibre Channel flow on Fabric Manager, follow these steps:

- Step 1** Choose **Switches > End Devices > SSM Features** from the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane, showing the FCWA tab.
- Step 2** Click the **Create Row** button in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow. You see the Fibre Channel write acceleration dialog box.
- Step 3** Select the initiator and target WWNs and VSAN IDs and check the **WriteAcc** check box to enable Fibre Channel write acceleration on this SCSI flow. You can optionally enable SCSI flow statistics on this SCSI flow at this time by checking the **Enable Statistics** check box.
- Step 4** Optionally, change the BufCount value to set the number of 2K buffers used by the SCSI target.
- Step 5** Click **Create** to create this SCSI flow or click **Cancel** to cancel this change.

## Fibre Channel Write Acceleration

Fibre Channel write acceleration minimizes application latency or reduces transactions per second over long distances. For synchronous data replication, Fibre Channel write acceleration increases the distance of replication or reduces effective latency to improve performance. To take advantage of this feature, both the initiator and target devices must be directly attached to an SSM.

The Fibre Channel write acceleration feature also allows the configuration of the buffer count. You can change the number of 2 KB buffers reserved on the target side for a SCSI flow.

You can estimate the number of buffers to configure using the following formula:

(Number of concurrent SCSI writes \* size of SCSI writes in bytes) / FCP data frame size in bytes



### Note

Fibre Channel write acceleration requires that the Enterprise Package license be installed on both the initiator and target switches.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Configuring Fibre Channel Write Acceleration

A SCSI flow is created when you configure Fibre Channel write acceleration. The SCSI flow consists of the following attributes:

- SCSI flow identifier
- VSAN identifier
- SCSI initiator port WWN
- SCSI target port WWN

To configure Fibre Channel write acceleration on Fabric Manager, and optionally modify the number of write acceleration buffers, follow these steps:

- 
- Step 1** Choose **Switches > End Devices > SSM Features** from the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane, showing the FCWA tab.
- Step 2** Click the **Create Row** button in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow. You see the Fibre Channel write acceleration dialog box.
- Step 3** Select the initiator and target WWNs and VSAN IDs and check the **WriteAcc** check box to enable Fibre Channel write acceleration on this SCSI flow. You can optionally enable SCSI flow statistics on this SCSI flow at this time by checking the **Enable Statistics** check box.
- Step 4** Optionally, set the BufCount value to set the number of 2K buffers used by the SCSI target.
- Step 5** Click **Create** to create this SCSI flow with Fibre Channel write acceleration or click **Cancel** to discard all changes.
- 

## SCSI Flow Statistics

The statistics that can be collected for SCSI flows include SCSI reads and writes, SCSI commands, and errors. To take advantage of this feature, only the initiator must be directly attached to an SSM.



### Note

SCSI flow statistics requires that the Enterprise Package license be installed only on the initiator switches.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Enabling SCSI Flow Statistics

To enable SCSI flow statistics monitoring using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; End Devices &gt; SSM Features</b> from the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane, showing the FCWA tab.  |
| <b>Step 2</b> | Click the <b>Create Row</b> button in the Information pane to create a SCSI flow or click a row in the FCWA table to modify an existing SCSI flow. You see the Fibre Channel write acceleration dialog box.   |
| <b>Step 3</b> | Select the the initiator and target WWNs and VSAN IDs and check the <b>Enable Statistics</b> check box to enable SCSI flow statistics on this SCSI flow. You can optionally enable Fibre Channel write acceleration on this SCSI flow at this time by checking the <b>WriteAcc</b> check box. |
| <b>Step 4</b> | Click <b>Create</b> to create this SCSI flow or click <b>Cancel</b> to cancel this change.  |
- 

## Viewing SCSI Flow Statistics and Clearing SCSI Flow Statistics

To clear SCSI flow statistics using Fabric Manager, follow these steps:


- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; End Devices &gt; SSM Features &gt; Statistics</b> from the Physical Attributes pane. You see the SCSI flow statistics displays in the Information pane. |
| <b>Step 2</b> | Choose <b>Switches &gt; End Devices &gt; SSM Features</b> from the Physical Attributes pane.  |
| <b>Step 3</b> | Click the <b>FCWA</b> tab in the Information pane. You see the FC write acceleration configuration in the Information pane.   |
| <b>Step 4</b> | Check the <b>Clear</b> check box in the statistics column to clear SCSI flow statistics.  |
| <b>Step 5</b> | Click the <b>Apply Changes</b> icon to clear the SCSI flow statistics or click the <b>Undo Changes</b> icon to discard any unsaved changes.                                     |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## SANTap

The SANTap feature allows third party data storage applications, such as long distance replication and continuous backup, to be integrated into the SAN. The protocol-based interface that is offered by SANTap allows easy and rapid integration of the data storage service application because it delivers a loose coupling between the application and an SSM, thereby reducing the effort needed to integrate applications with the core services being offered by the SSM. See [Figure 22-1](#).

**Figure 22-1 Integrating Third-Party Storage Applications in a SAN**



SANTap operates in three modes:


- Transparent mode
- Proxy mode-1
- Proxy mode-2

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Transparent Mode

Transparent mode eliminates the need for any reconfiguration of either the host or target when introducing SANTap based applications. This mode of operation requires that either the host initiator or target is directly connected to an SSM. See [Figure 22-2](#).

**Figure 22-2** *SANTap Transparent Mode Example*



130076




***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Proxy Mode-1

Proxy mode-1 assigns Cisco-specific WWNs to the virtual initiators (VIs) and digital virtual targets (DVTs). The benefit of this mode is that it eliminates the transparent mode requirement that a host initiator or a target be connected directly to an SSM. In proxy mode-1, the SSM can be anywhere in the SAN. However, this mode requires reconfiguration of legacy applications. See [Figure 22-3](#).

**Figure 22-3** *SANTap Proxy Mode-1 Example*




**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Proxy Mode-2

Proxy mode-2 includes the benefits of transparent mode and proxy mode-1 but does not have the limitations of those modes. However, it does require that the administrator partition the SAN using VSANs. The host initiator and the DVT are in one VSAN while the VI and the target are in another VSAN. See [Figure 22-4](#).

**Figure 22-4 SANTap Proxy Mode-2 Example**



## Configuring SANTap

To configure SANTap using Fabric Manager, follow these steps:

- Step 1** Choose **Switches > End Devices > SSM Features** from the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane, showing the FCWA tab.
- Step 2** Click the **SANTap** tab. You see the SANTap configuration in the Information pane.
- Step 3** Click the **Create Row...** icon. You see the SANTap configuration dialog box.
- Step 4** Select the **Switch** and **SSM Card** you want to configure SANTap on.



**Note** SANTap must be enabled and provisioned as a service on this SSM module. See the [“Enabling Intelligent Storage Services”](#) section on page 22-2.

- Step 5** Select the VSAN ID you want to configure for SANTap.
- Step 6** Click **Create** to create this SANTap or click **Cancel** to cancel this change.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## NASB

As of Cisco MDS SAN-OS Release 2.1(1), SSMs support Network-Accelerated Serverless Backup (NASB).

Data movement in the fabric uses considerable processor cycles that can cause client applications to slow down noticeably. Offloading data movement operations to a media server allows the client applications to run normally even during a backup operation. Media servers can further offload the data movement operation to NASB devices, which allows the media server to focus on the coordination functions needed to complete the backup.

Most backups performed today are server-free. In server-free backups, the application server is not involved in moving the data. The data can be moved by either a media server or an NASB device.

When the media server is the data mover, it moves the data between the disks and the tapes. The backup application runs on both the client device and the media server. However, the backup application in the client device performs minimal tasks for the backup operation.


The media server performs the following backup operations:

- Manages disks as well as one or more tape backup devices.
- Contacts the client devices to retrieve the list of logical blocks that need to be backed up.
- Performs data movement from disk to tape media based on the logical block list provided by the client device.

The backup application in the client device maps the data to be backed up and creates the logical block list associated with the data. The movement of data from the physical disks to the backup device (tape) is not performed by the client device. This reduces substantial load on the client device.

An example configuration is shown in [Figure 22-5](#). The media server moves the data directly between the storage disks and the tape devices during backups.

**Figure 22-5 Example Configuration with Media Server as Data Mover**

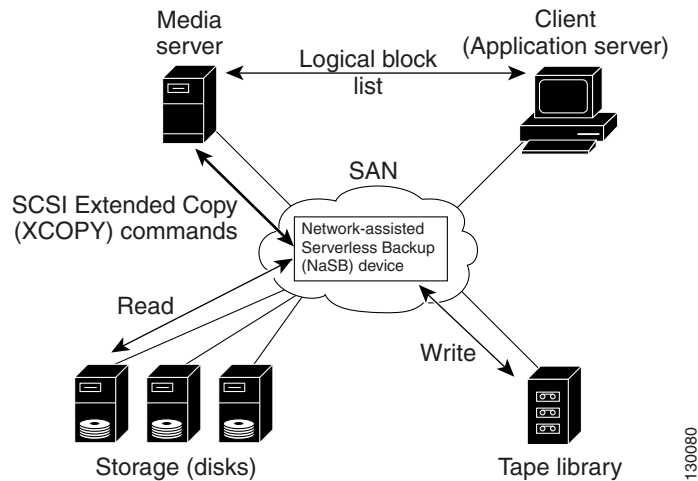


130079

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

When the NASB is the data mover, it moves the data between the disks and the tapes. The NASB device is a SCSI target device capable of handling SCSI Extended Copy (XCOPY) commands as well as a SCSI initiator device capable of issuing read/write commands to disks and other backup media, such as tapes. See Figure 22-6.

**Figure 22-6 Example Configuration with NASB Device as Data Mover**



The task of managing and preparing the source and destination targets is performed by the media server. For example, if the destination is a tape library, the media server issues commands to load and unload the correct tape and position of the tape write head at the correct offset within the tape.

## Configuring NASB

VSANs used with NASB must be configured to permit default zoning.

To permit default zoning on a VSAN using Fabric Manager, follow these steps:

- 
- Step 1** Choose **VSANx.xx > Default Zone**. You see the Default Zone configuration in the Information pane.
  - Step 2** Choose the **Policies** tab. You see the default zone policy for the selected VSAN.
  - Step 3** Choose permit from the **Default Zone > Behavior** drop-down box to permit default zoning on this VSAN.
  - Step 4** Click the **Apply Changes** icon to clear the SCSI flow statistics or click the **Undo Changes** icon to discard any unsaved changes.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To configure NASB using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > End Devices > SSM Features** from the Physical Attributes pane. You see the Intelligent Storage Services configuration in the Information pane, showing the FCWA tab.
- Step 2** Click the **NASB** tab. You see the NASB configuration in the Information pane.
- Step 3** Click the **Create Row...** icon. You see the NASB configuration dialog box.
- Step 4** Select the **Switch** and **SSM Card** you want to configure NASB on.



---

**Note** The NASB must be enabled and provisioned as a service on this SSM module. See the [“Enabling Intelligent Storage Services”](#) section on page 22-2.

---

- Step 5** Select the VSAN ID you want to configure for NASB.



---

**Note** You must configure this VSAN to permit default zoning.

---

- Step 6** Click **Create** to create this NASB or click **Cancel** to cancel this change.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Additional Configuration

This chapter describes the advanced features provided in switches in the Cisco MDS 9000 Family.



**Note**

Click **Help > Contents** in Fabric Manager to access help topics not covered in this document.

This chapter includes the following sections:

- [Fibre Channel Time Out Values, page 23-1](#)
- [Configuring World Wide Names, page 23-3](#)
- [Flat FC ID Allocation, page 23-4](#)
- [Switch Interoperability, page 23-4](#)

## Fibre Channel Time Out Values

To configure timers in Fabric Manager, choose **Switches > FC Services > Timers & Policies** on the Physical Attributes tree. You see the timers for multiple switches in the Information pane. Click the **Change Timeouts** button to configure the timeout values.

To configure timers in Device Manager, choose **FC > Advanced > Timers/Policies**. You see timers for a single switch in the dialog box.

You can modify Fibre Channel protocol related timer values for the switch by configuring the following time out values (TOVs):

- Distributed services TOV (D\_S\_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 5,000 milliseconds.
- Error detect TOV (E\_D\_TOV)—The valid range is from 1,000 to 10,000 milliseconds. The default is 2,000 milliseconds. This value is matched with the other end during port initialization.
- Resource allocation TOV (R\_A\_TOV)—The valid range is from 5,000 to 10,000 milliseconds. The default is 10,000 milliseconds. This value is matched with the other end during port initialization.



**Note**

The fabric stability TOV (F\_S\_TOV) constant cannot be configured.

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

## The fctrace Feature

The fctrace feature allows you to:

- Trace the route followed by data traffic.
- Compute inter-switch (hop-to-hop) latency.

You can invoke fctrace by providing the FC ID, the N port, or the NL port WWN of the destination. The frames are routed normally as long as they are forwarded through TE ports.

Once the frame reaches the edge of the fabric (the F port or FL port connected to the end node with the given port WWN or the FC ID), the frame is looped back (swapping the source ID and the destination ID) to the originator.

If the destination cannot be reached, the path discovery starts and traces the path up to the point of failure.



### Note

The fctrace feature works only on TE ports. Make sure that only TE ports exist in the path to the destination. In case there is an E port in the path, the fctrace frame is dropped by that switch. Also, fctrace times out in the originator, and path discovery does not start.



### Tip

You cannot use the fctrace feature in a locally configured VSAN interface (IPFC interface), but you can trace the route to a VSAN interface configured in other switches.

## Performing an fctrace Operation

To initiate an fctrace using Fabric Manager, follow these steps:

- Step 1** Choose **Tools > Traceroute**. You see the fctrace dialog box.
- Step 2** Set the Source Switch, VSAN, and Target Endport for where you want the trace to begin and end.
- Step 3** Click **Start** to start the fctrace. You see updates to the dialog box for each hop. This shows each switch the trace traverses from the source switch to the target end point. If the destination is unreachable, you see the hops that the trace traversed until the point where the target could not be found.
- Step 4** Click **Close** to close the dialog box.

## The fcping Feature

The fcping feature verifies reachability of a node by checking its end-to-end connectivity. You can invoke the fcping feature by providing the FC ID or the destination port WWN information.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Invoking the fcping Feature

To initiate an fcping using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Tools &gt; Ping</b> . You see the ping dialog box.                                  |
| <b>Step 2</b> | Set the Source Switch, VSAN, and Target Endport for where you want the ping to begin and end. |
| <b>Step 3</b> | Click <b>Start</b> to start the fcping. You see the results of the ping in the dialog box.    |
| <b>Step 4</b> | Click <b>Close</b> to close the dialog box.   |
- 

## Configuring World Wide Names

The world wide name (WWN) in the switch is equivalent to the Ethernet MAC address. As with the MAC address, you must uniquely associate the WWN to a single device. The principal switch selection and the allocation of domain IDs rely on the WWN. The WWN manager, a process-level manager residing on the switch's supervisor module, assigns WWNs to each switch.

To access information on the WWN from Fabric Manager, choose **Switches > FC Services > WWN Manager**.

Cisco MDS 9000 Family switches support three network address authority (NAA) address formats (see [Table 23-1](#)).

**Table 23-1**      **Standardized NAA WWN Formats**

NAA Address	NAA Type	WWN Format	
IEEE 48-bit address	Type 1 = 0001b	000 0000 0000b	48-bit MAC address
IEEE extended	Type 2 = 0010b	Locally assigned	48-bit MAC address
IEEE registered	Type 5 = 0101b	IEEE company ID: 24 bits	VSID: 36 bits



### Caution

Changes to the world-wide names should be made by an administrator or individual who is completely familiar with switch operations.

## Link Initialization WWN Usage

Exchange Link Protocol (ELP) and Exchange Fabric Protocol (EFP) use WWNs during link initialization. The usage details differ based on the Cisco MDS SAN-OS software release:

- In Cisco MDS SAN-OS Release 1.0 and 1.1, both ELPs and EFPs use the VSAN WWN during link initialization.
- In Cisco MDS SAN-OS Releases 1.2 and 1.3, two different WWNs are used during the link initialization process:
  - ELPs use the switch WWN.
  - EFPs use the VSAN WWN.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- In Cisco MDS SAN-OS Release 2.0, both ELPs and EFPs use the VSAN WWN by default during link initialization. However, the ELP usage changes based on the peer switch's usage:
  - If the peer switch ELP uses the switch WWN, then the local switch also uses the switch WWN.
  - If the peer switch ELP uses the VSAN WWN, then the local switch also uses the VSAN WWN.

This link initialization change between Cisco MDS SAN-OS releases is implicit and does not require any configuration.

## Flat FC ID Allocation

Fibre Channel standards require a unique FC ID to be allocated to an N port attached to a Fx port in any switch. To conserve the number of FC IDs used, Cisco MDS 9000 Family switches use a special allocation scheme.

Based on the assigned FC ID, some HBAs assume that no other ports have the same area bits and domain. When a target is assigned with an FC ID that has the same area bits, but different port bits, the HBA fails to discover these targets. To isolate these HBAs in a separate area, switches in the Cisco MDS 9000 Family follow a different FC ID allocation scheme. By default, the FC ID allocation mode is auto mode. In the auto mode, only HBAs without interop issues are assigned FCIDs with specific port bits. All other HBAs are assigned FC IDs with a whole area (port bits set to 0). The three options to allocate FCID are auto (default), none, and flat.



**Caution**

Changes to FC IDs should be made by an administrator or individual who is completely familiar with switch operations.

## Loop Monitoring Initiation

By default, the loop monitoring is disabled in all switches in the Cisco MDS 9000 Family. When a disk is removed from a loop port, the loop stays active based on the bypass circuit. Thus the disk removal is not known until you try to communicate with the disk. To detect such removals, the disks can be polled periodically (every 20 seconds).



**Caution**

Changes to the loop monitoring feature should be made by an administrator or individual who is completely familiar with switch operations.

## Switch Interoperability

Interoperability enables the products of multiple vendors to come into contact with each other. Fibre Channel standards guide vendors towards common external Fibre Channel interfaces.

If all vendors followed the standards in the same manner, then interconnecting different products would become a trivial exercise. However, not all vendors follow the standards in the same way thus resulting in interoperability modes. This section briefly explains the basic concepts of these modes.

Each vendor has a regular mode and an equivalent interoperability mode, which specifically turns off advanced or proprietary features and provides the product with a more interoperable standards compliant implementation.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Table 23-2 lists the changes in switch behavior when you enable interoperability mode. These changes are specific to switches in the Cisco MDS 9000 Family while in interop mode.

**Table 23-2 Changes in Switch Behavior When Interoperability Is Enabled**

Switch Feature	Changes if Interoperability Is Enabled
Domain IDs	Some vendors cannot use the full range of 239 domains within a fabric.  Domain IDs are restricted to the range 97-127. This is to accommodate McData's nominal restriction to this same range. They can either be set up statically (the Cisco MDS switch accept only one domain ID, if it does not get that domain ID it isolates itself from the fabric) or preferred. (If it does not get its requested domain ID, it accepts any assigned domain ID.)
Timers	All Fibre Channel timers must be the same on all switches as these values are exchanged by E ports when establishing an ISL. The timers are F_S_TOV, D_S_TOV, E_D_TOV, and R_A_TOV.
F_S_TOV	Verify that the Fabric Stability Time Out Value timers match exactly.
D_S_TOV	Verify that the Distributed Services Time Out Value timers match exactly.
E_D_TOV	Verify that the Error Detect Time Out Value timers match exactly.
R_A_TOV	Verify that the Resource Allocation Time Out Value timers match exactly.
Trunking	Trunking is not supported between two different vendor's switches. This feature may be disabled on a per port or per switch basis.
Default zone	The default zone behavior of permit (all nodes can see all other nodes) or deny (all nodes are isolated when not explicitly placed in a zone) may change.
Zoning attributes	Zones may be limited to the pWWN and other proprietary zoning methods (physical port number) may be eliminated.  <b>Note</b> Brocade uses the <b>cfgsave</b> command to save fabric-wide zoning configuration. This command does not have any effect on Cisco MDS 9000 Family switches if they are part of the same fabric. You must explicitly save the configuration on each switch in the Cisco MDS 9000 Family.
Zone propagation	Some vendors do not pass the full zone configuration to other switches, only the active zone set gets passed.  Verify that the active zone set or zone configuration has correctly propagated to the other switches in the fabric.
VSAN	Interop mode only affects the specified VSAN.
TE ports and PortChannels	TE ports and PortChannels cannot be used to connect Cisco MDS to non-Cisco MDS switches. Only E ports can be used to connect to non-Cisco MDS switches. TE ports and PortChannels can still be used to connect an Cisco MDS to other Cisco MDS switches even when in interop mode.
FSPF	The routing of frames within the fabric is not changed by the introduction of interop mode. The switch continues to use src-id, dst-id, and ox-id to load balance across multiple ISL links.
Domain reconfiguration disruptive	This is a switch-wide impacting event. Brocade and McData require the entire switch to be placed in offline mode and/or rebooted when changing domain IDs.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 23-2**      **Changes in Switch Behavior When Interoperability Is Enabled (continued)**

Switch Feature	Changes if Interoperability Is Enabled
Domain reconfiguration nondisruptive	This event is limited to the affected VSAN. Only Cisco MDS 9000 Family switches have this capability—only the domain manager process for the affected VSAN is restarted and not the entire switch.
Name server	Verify that all vendors have the correct values in their respective name server database.
IVR	IVR-enabled VSANs can only be configured in either the non interop mode or interop mode 1.

## Interoperability Configuration

The interop mode in Cisco MDS 9000 Family switches can be enabled disruptively or nondisruptively.



### Note

Brocade's `msplmgmtdeactivate` command must explicitly be run prior to connecting from a Brocade switch to either Cisco MDS 9000 Family switches or to McData switches. This command uses Brocade proprietary frames to exchange platform information, which Cisco MDS 9000 Family switches and McData switches do not understand. Rejecting these frames causes the common E ports to become isolated.

## Configuring Interoperability

To configure the interop mode for a VSAN using Fabric Manager, follow these steps:

- Step 1** Choose **VSANxxx > VSAN Attributes** from the Logical Domains pane. You see the VSAN attributes in the Information pane.
- Step 2** Select **Interop-1** in the InterOp drop-down box.
- Step 3** Choose **Apply Changes** to save this interop mode.
- Step 4** Choose **VSANxxx > Domain Manager** from the Logical Domains pane. You see the Domain Manager configuration in the Information pane.
- Step 5** Set the domain ID in the range of 97 (0x61) through 127 (0x7F).



### Note

This is a limitation imposed by the McData switches.



### Note

When changing the domain ID, the FC IDs assigned to N ports also change.

- Step 6** Change the Fibre Channel timers (if they have been changed from the system defaults).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

The Cisco MDS 9000, Brocade, and McData FC Error Detect (ED\_TOV) and Resource Allocation (RA\_TOV) timers default to the same values. They can be changed if needed. The RA\_TOV default is 10 seconds, and the ED\_TOV default is 2 seconds. Per the FC-SW2 standard, these values must be the same on each switch within the fabric.

- a. Choose **Switches > FC Services > Timers and Policies**. You see the timer settings in the Information pane.
- b. Click **Change Timeouts** to modify the timeout values.
- c. Click **Apply** to save the new timeout values.

**Step 7** Optionally, choose **VSANxxx > Domain Manager** and select **disruptive** or **nonDisruptive** in the restart drop-down box to restart the domain.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 4**

### **Security Configuration**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Users and Common Roles

---

Fabric Manager provides the capability to configure and manage several different types of security for MDS 9000 switches.

This chapter includes the following sections:

- [Role-Based Authorization, page 24-1](#)
- [Configuring Common Roles, page 24-2](#)
- [Configuring User Accounts, page 24-4](#)
- [Configuring SSH Services, page 24-6](#)

## Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context sensitive help, the switch software allows the operation to progress if you have permission to access that command.

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to debug commands, then if Joe belongs to both role1 and role2, he can access configuration as well as debug commands.



### Note

If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Roles are cumulative. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.



### Tip

Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Configuring Common Roles

From Release 1.2(x), CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles.

You can use SNMP to modify a role that was created using CLI and vice versa. Each role in SNMP is the same as a role created or modified through the CLI. Common Roles allow you to use a set of rules to set the scope of VSAN security. Each role can be restricted to one or more VSANs as required.

To configure Common Roles from the Device Manager, select Common Roles from the Security menu. You can then access the Rules dialog box to configure the set of rules.

To configure Common Roles from Fabric Manager, select Security > SNMP and click the Roles tab in the Information pane. Fabric Manager uses a default rules set for roles; therefore, no Rules dialog box is displayed.



### Note

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is clicked, the other Information pane tabs that use CFS are activated.

## Creating Common Roles

To create a common role, follow these steps.

- Step 1** In Fabric Manager, choose **Switches > Security > SNMP** from the Physical Attributes , and click the **Roles** tab in the Information pane.  
In Device Manager, choose **Common Roles** from the Security menu. You see the Common Roles dialog box.
- Step 2** Click the **Create Row** icon to create a new role in Fabric Manager or click **Create** in Device Manager. You see the Roles - Create dialog box.
- Step 3** Select the switches on which you want to configure the role in Fabric Manager.
- Step 4** Enter the name of the role in the Name field.
- Step 5** Enter the description of the role in the Description field.
- Step 6** Check the **Has Config and Exec Permission** check box if you want your role to have read, write, and create permission. If you do not check the **Has Config and Exec Permission** check box, your role will have read-only permission.
- Step 7** Optionally, check the **Enable** check box to enable the VSAN scope and enter the list of VSANs in the Scope field that you want to restrict this role to.
- Step 8** Click **Create** to create the role, or click **Close** to close the Roles - Create dialog box without creating the common role.



### Note

Device Manager automatically creates six roles that are required for Device Manager to display a view of a switch. These roles are: **system**, **snmp**, **module**, **interface**, **hardware**, and **environment**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Editing Rules For Common Roles in Device Manager

Up to 16 rules can be configured for each role. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

**Note**

The order of rule placement is important. If you place a more permissive policy after a restrictive policy, the permissive policy may have priority over the restrictive policy.

To edit the rules for a common role in Device Manager, follow these steps.

- 
- Step 1** Choose **Security > Roles**. You see the Common Roles dialog box.
  - Step 2** Click the common role that you want to edit the rules for.
  - Step 3** Click **Rules** to view the rules for the role. You see the Rules dialog box. It may take a few minutes to display.
  - Step 4** Edit the rules you want to enable or disable for the common role.
  - Step 5** Click **Apply** to apply the new rules and close the Rules dialog, or click **Close** to close the Rules dialog without applying the rules.
- 

## Deleting Common Roles

To delete a common role, follow these steps:

- 
- Step 1** In Fabric Manager, choose **Switches > Security > SNMP** from the Physical Attributes pane and click the **Roles** tab in the Information pane.  
In Device Manager, choose **Security > Common Roles**. You see the Common Roles dialog box.
  - Step 2** Click the common role you want to delete.
  - Step 3** Click the **Delete Row** icon in Fabric Manager or **Delete** in Device Manager to delete the common role.
- 

## Configuring the VSAN Policy

Configuring the VSAN policy or VSAN scope requires the ENTERPRISE\_PKG license. See [Chapter 9, “Obtaining and Installing Licenses.”](#)

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN scope for any role is disabled. That is, the roles allow tasks to be performed in all VSANs. To configure a role to selectively allow tasks in a subset of VSANs, you must enable the VSAN scope and then list the appropriate VSANs in the VSAN list.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

Users configured in roles where the VSAN scope is enabled cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.

**Tip**

Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN scope is enabled are referred to as VSAN-restricted users. These users cannot perform tasks that require the startup configuration to be viewed or modified.

## Modifying the VSAN Policy

To modify the VSAN policy or VSAN scope for an existing common role, follow these steps.

- Step 1** In Fabric Manager, choose **Switches > Security > SNMP** from the Physical Attributes , and click the **Roles** tab in the Information pane.  
In Device Manager, choose **Common Roles** from the Security menu. You see the Common Roles dialog box.
- Step 2** Check the **enable** check box if you want to enable the VSAN scope and restrict this role to a subset of VSANs.
- Step 3** Enter the list of VSANs in the VSAN Scope > List field that you want to restrict this role to.
- Step 4** Click **Apply Changes** in Fabric Manager or click **Apply** in Device Manager to save these changes. Click **Undo Changes** in Fabric Manager or click **Close** in Device Manager to discard any unsaved changes.

## Configuring User Accounts

Every Cisco MDS 9000 Family switch user has the account information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile.

The tasks explained in this section enable you to create users and modify the profile of an existing user. These tasks are restricted to privileged users as determined by your administrator.

**Note**

Cisco SAN-OS does not support all numeric usernames, whether created with TACACS+ or RADIUS, or created locally. Local users with all numeric names cannot be created. If an all numeric username exists on an AAA server and is entered during login, the user is not logged in.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Creating or Updating Users

As of Cisco MDS SAN-OS Release 2.x and later, the passphrase specified in the **snmp-server user** option and the password specified in the **username** option are synchronized.

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** CLI option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.



### Tip

The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nsd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.



### Note

User passwords are not displayed in the switch configuration file.

## Creating Strong Passwords

If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. As of Cisco MDS SAN-OS Release 2.x and later, admin is not the default password for any switch in the Cisco MDS 9000 Family. You must explicitly configure a password that meets the following requirements:

- Is at least eight characters in length.
- Does not have multiple consecutive characters (such as abcd or jklm).
- Does not have multiple repeat characters (such as fff or qqdd).
- Does not contain words found in a dictionary.

## Adding a User

To add a user, follow these steps:

- Step 1** In Fabric Manager, choose **Switches > Security > SNMP** from the Physical Attributes pane and click the **Users** tab in the Information pane.  
In Device Manager, choose **Security > SNMP** and click the **Users** tab.
- Step 2** Click the **Create Row** icon in Fabric Manager or click **Create** in Device Manager. You see the Create Users dialog box.  
The dialog box from Fabric Manager also provides check boxes to specify one or more switches.
- Step 3** Enter the user name in the **New User** field.
- Step 4** Select the role from the drop-down menu in Fabric Manager or the check boxes in Device Manager. In Fabric Manager, you can enter a new role name in the field if you do not want to select one from the drop-down menu. If you do this, you must go back and configure this role appropriately (see the [“Configuring Common Roles” section on page 24-2](#)).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Enter the password for the user twice in the New Password and Confirm Password fields.
  - Step 6** Check the **Privacy** check box and complete the password fields to encrypt management traffic.  
Enter the same new password in the New Password and Confirm Password fields.
  - Step 7** Click **Create** to create the new entry or click **Close** to discard any unsaved changes and close the dialog box.
- 

## Deleting a User

To delete a user, follow these steps:

- Step 1** In Fabric Manager, choose **Switches > Security > SNMP** from the Physical Attributes pane and click the **Users** tab in the Information pane.  
In Device Manager, choose **Security > SNMP** and click the **Users** tab.
  - Step 2** Click the name of the user you want to delete.
  - Step 3** Click the **Delete Row** icon in Fabric Manager or **Delete** in Device Manager to delete the selected user.
- 

## Viewing User Information

To view information about users, follow these steps:

- Step 1** In Fabric Manager, choose **Security > SNMP** from the Physical Attributes pane.  
In Device Manager, choose **Security > SNMP**. You see the
  - Step 2** Click the **Users** tab in Fabric Manager. You see the list of SNMP users in the Information pane.
- 

## Configuring SSH Services

The Telnet service is enabled by default on all Cisco MDS 9000 Family switches. Before enabling the SSH service, generate a server key pair.

## Generating the SSH Server Key Pair and Enabling SSH

Be sure to have an SSH server key pair with the appropriate version before enabling the SSH service. Generate the SSH server key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The SSH service accepts three types of key pairs for use by SSH versions 1 and 2.

- TheSSH1 option generates the RSA1 key pair for the SSH version 1 protocol.
- TheSSH2(dsa) option generates the DSA key pair for the SSH version 2 protocol.
- The SSH2(rsa) option generates the RSA key pair for the SSH version 2 protocol.



**Caution** If you delete all of the SSH keys, you cannot start a new SSH session.

To generate an SSH server key pair and enable SSH, follow these steps:

- 
- Step 1** In Fabric Manager, choose **Switches > Security > SSH**. You see the SSH configuration in the Information pane.  
In Device Manager, choose **Security > SSH**. You see the SSH dialog box.
- Step 2** Click the **Create Row** icon in Fabric Manager or click **Create** in Device Manager. You see the SSH Key Create dialog box.
- Step 3** Optionally, check the switches you want this SSH key pair for in Fabric Manager.
- Step 4** Choose the **Control** tab in Fabric Manager check the **enable** check box to enable SSH on the selected switches.  
Check the **enable** check box in Device Manager to enable SSH.
- Step 5** Choose the key pair option type from the Protocols radio buttons.
- Step 6** Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.
- Step 7** Click **Create** to generate these keys or click **Close** to discard any unsaved changes.
- 

## Deleting a Generated Key Pair

If the SSH key pair option is already generated for the required SSH protocol version, you must delete the previously generated key pair before you can a new key pair for that SSH protocol version.

## Recovering Administrator Password

An administrator can recover a password from a local console connection. The password recovery procedure must be performed on the supervisor module that becomes the active supervisor module after the recovery procedure is completed.



**Note** To recover a n administrator's password, refer to the *Cisco MDS 9000 Family Configuration Guide*.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## SNMP Configuration

Fabric Manager provides the capability to configure SNMP for managing switches in the fabric.

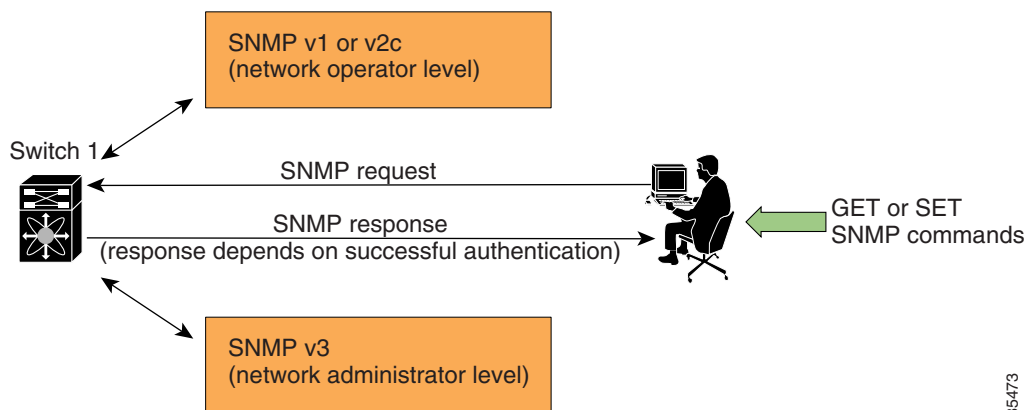
This chapter includes the following sections:

- [About SNMP, page 25-1](#)
- [Understanding Users, page 25-4](#)
- [Assigning SNMPv3 Users to Multiple Roles, page 25-5](#)
- [Configuring SNMP Notifications, page 25-6](#)

### About SNMP

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see [Figure 25-1](#)).

**Figure 25-1**      **SNMP**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## SNMP Version 1 and Version 2c

SNMPv1 and SNMPv2c use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

## SNMP Version 3

SNMPv3 is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources. Uses DES or AES.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

## SNMP v3 CLI User Management and AAA Integration

The Cisco MDS SAN-OS software implement RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

As of Cisco MDS SAN-OS Release 2.0(1b), SNMP v3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of AAA server. Once user authentication is verified, the SNMP PDUs are processed further. Additionally, the AAA server is also used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

## CLI and SNMP User Synchronization

In Cisco MDS SAN-OS Release 2.0(1b) or later, all updates to the CLI security database and the SNMP user database are synchronized. You can use the CLI password for accessing Fabric Manager or Device Manager and CLI. After you upgrade to Cisco MDS SAN-OS Release 2.0(1b) or later, you can continue using the SNMP password for Fabric Manager or Device Manager. If you use the CLI password for Fabric Manager or Device Manager login, you need to use the CLI password for future logins as well.

In Cisco MDS SAN-OS Release 2.0(1b) or later, users present in the prior release are assigned set of roles that is the union of both the CLI and the SNMP rules. Any configuration changes made to the user group, role, or password, results in the database synchronization for both SNMP and AAA.



### Note

When the passphrase/password is specified in localized key/encrypted format, the password is not synchronized.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Software Upgrade Synchronization

When you upgrade from an earlier release to Cisco MDS SAN-OS Release 2.0(1b) or later, the following synchronization steps occur:

- Existing SNMP users continue to retain the `auth` and `priv` information without any changes.
- If a user is not present in one database and is present in other database, the CLI user is created without any password (login is disabled) and the SNMP user is created with the `noAuthNoPriv` security level. Subsequently, the passwords and roles for these users will be synchronized.
- If the management station creates a SNMP user in the `usmUserTable`, this user is created without any password (login is disabled) and will have the `network-operator` role.

## Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP Access Control Lists (IP-ACLs). See the [“IP-ACL Configuration Guidelines”](#) section on page 27-1.

## Adding a Community String

To add a community string, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | From Fabric Manager, choose <b>Switches &gt; Security-&gt;SNMP</b> from the Physical Attributes pane and click the <b>Communities</b> tab in the Information pane.<br><br>From Device Manager, choose <b>Security &gt; SNMP</b> and click the <b>Communities</b> tab.   |
| <b>Step 2</b> | Click <b>Create</b> in the Device Manager dialog box, or click the <b>Create Row</b> icon in Fabric Manager .<br><br>You see the Create Community string dialog box.<br><br>The dialog box in Fabric Manager also provides check boxes to specify one or more switches.   |
| <b>Step 3</b> | Enter the community name in the Community field.  |
| <b>Step 4</b> | Select the role from the check boxes in Device Manager or the drop-down list in Fabric Manager. In Fabric Manager, you can enter a new role name in the field if you do not want to select one from the drop-down list. If you do this, you must go back and configure this role appropriately (see the <a href="#">“Configuring Common Roles”</a> section on page 24-2). |
| <b>Step 5</b> | Click <b>Create</b> to create the new entry or click <b>Close</b> to create the entry and close the dialog box.   |
- 

## Deleting a Community String

To delete a community string, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | From Fabric Manager, choose <b>Switches &gt; Security-&gt;SNMP</b> from the Physical Attributes pane and click the <b>Communities</b> tab in the Information pane.<br><br>From Device Manager, choose <b>Security &gt; SNMP</b> and click the <b>Communities</b> tab. |
| <b>Step 2</b> | Click the name of the community you want to delete.   |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 3** Click **Delete** in Device Manager or click the **Delete Row** icon in Fabric Manager.
- 

## Understanding Users

Every Cisco MDS 9000 Family switch user has the account information stored by the system. Your authentication information, user name, user password, password expiration date, and role membership are stored in your user profile.

## Adding a User

To add a user, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Security->SNMP** from the Physical Attributes pane and click the **Communities** tab in the Information pane.
- From Device Manager, choose **Security > SNMP** and click the **Users** tab.
- Step 2** Click the **Create Row** icon in Fabric Manager or click **Create** in the Device Manager dialog box.
- You see the Create Users dialog box.
- The dialog box from Fabric Manager also provides check boxes to specify one or more switches.
- Step 3** Enter the user name in the New User field.
- Step 4** Select the role from the check boxes in Device Manager or the drop-down list in Fabric Manager. In Fabric Manager, you can enter a new role name in the field if you do not want to select one from the drop-down list. If you do this, you must go back and configure this role appropriately (see the [“Configuring Common Roles”](#) section on page 24-2).
- Step 5** Enter the same authentication password for the user in the New Password and Confirm Password fields.
- Step 6** Check the **Privacy** check box and complete the password fields to enable encryption of management traffic.
- Enter the same new privacy password in the New Password and Confirm Password fields.
- Step 7** Click **Create** to create the new entry or click **Close** close the dialog box without creating an entry.
- 

## Deleting a User

To delete a user, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Security->SNMP** from the Physical Attributes pane and click the **Communities** tab in the Information pane.
- From Device Manager, choose **Security > SNMP** and click the **Users** tab.
- Step 2** Click the name of the user you want to delete.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 3** Click **Delete** in Device Manager or click the **Delete Row** icon in Fabric Manager.

---

## Viewing SNMP Community and User Information

To view information about SNMP users, roles, and communities from Fabric Manager, choose **Security > SNMP** from the Physical tree and click the **Users**, **Roles**, or **Communities** tab. You see the list of SNMP users, roles, or communities in the Information pane.

To view this information from the Device Manager, choose **SNMP** from the Security menu. You see the SNMP dialog box.

## Group-Based SNMP Access



### Note

Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

## Assigning SNMPv3 Users to Multiple Roles

As of Cisco MDS SAN-OS Release 2.0(1b), the SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. You map additional roles for the user at the time you create the user.



### Note

Only users belonging to network-admin role can assign roles to other users.

To add multiple roles to a new user using Device Manager, follow these steps:

- 
- Step 1** Choose **Security > SNMP** and click the **Users** tab.
  - Step 2** Click **Create** in the Device Manager dialog box. You see the Create Communities dialog box.
  - Step 3** Enter the user name in the New User field.
  - Step 4** Select the multiple roles from the check boxes in Device Manager .
  - Step 5** Enter the same authentication password for the user in the New Password and Confirm Password fields.
  - Step 6** Check the **Privacy** check box and complete the password fields to enable encryption of management traffic.  
Enter the same new privacy password in the New Password and Confirm Password fields.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

- Step 7** Click **Create** to create the new entry or click **Close** close the dialog box without creating an entry.
- 

## Configuring SNMP Notifications

You can configure the Cisco MDS switch to send notifications to SNMP managers when particular events occur. You can send these notifications as traps or inform requests. Traps are unreliable because the receiver does not send any acknowledgment when it receives a trap. The sender cannot determine if the trap was received. However, an SNMP manager that receives an inform request acknowledges the message with an SNMP response PDU. If the sender never receives a response, the inform request can be sent again. Thus, informs are more likely to reach their intended destination.



### Note

Use the SNMP-TARGET-MIB to obtain more information on trap destinations and inform requests. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information.

---

To configure SNMP notifications (traps or informs) using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Events > SNMP Trap** in the Physical Attributes pane. You see the SNMP notification configuration in the Information pane.
  - Step 2** Click the **Destinations** tab to add or modify a receiver for SNMP notifications.
  - Step 3** Click **Create Row** to create a new notification destination. You see the Create Destination Dialog box.
  - Step 4** Check the switches that you want to configure a new destination on.
  - Step 5** Set the destination IP address and UDP port.
  - Step 6** Choose either the **trap** or **inform** radio button.
  - Step 7** Optionally, set the inform timeout and retry values.
  - Step 8** Click **Create** to add this destination to the selected switches or click **Close** to discard any unsaved changes.
  - Step 9** Optionally, click the other tabs to enable specific notification types per switch.
  - Step 10** Click the **Apply changes** icon to create the entry or click **Undo Changes** to discard any unsaved changes.
-



## RADIUS and TACACS+

---

Fabric Manager provides the capability to authenticate users with RADIUS or TACACS+.

This chapter includes the following sections:

- [Authentication, Authorization, and Accounting, page 26-1](#)
- [Configuring RADIUS, page 26-5](#)
- [Configuring TACACS+, page 26-7](#)
- [Configuring Server Groups, page 26-9](#)

## Authentication, Authorization, and Accounting

The authentication, authorization, and accounting (AAA) mechanism verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use Remote Access Dial-In User Service (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+) protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using AAA server(s). A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA server or for only a specific AAA server. This security mechanism provides a central management capability for AAA servers.

## CLI Security Options

You can access the CLI using the console (serial connection), Telnet, or Secure Shell (SSH). For each management path (console or Telnet and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
  - Using Remote Authentication Dial-In User Services (RADIUS). See the [“Configuring RADIUS” section on page 26-5](#).
  - Using Terminal Access Controller Access Control System plus (TACACS+). See the [“Configuring TACACS+” section on page 26-7](#).
- Local security control. See the [“Local AAA Services” section on page 26-11](#).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

These security mechanisms can also be configured for the following scenarios:

- iSCSI authentication (see the “[iSCSI User Authentication](#)” section on page 20-17).
- Fibre Channel Security Protocol (FC-SP) authentication (see the “[Fibre Channel Security Protocol](#)” section on page 29-1)

## SNMP Security Options

The SNMP agent supports security features for SNMPv1, SNMPv 2c, and SNMPv3. Normal SNMP security mechanisms apply to all applications that use SNMP (for example, Cisco MDS 9000 Fabric Manager).

Fabric Manager and Device Manager security options also apply to the CLI.

See the “[SNMP Version 3](#)” section on page 25-2.

## Switch AAA Functionalities

Using Fabric Manager, you can configure authentication, authorization, and accounting (AAA) switch functionalities on any switch in the Cisco MDS 9000 Family.

## Authentication

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person trying to manage the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).



### Note

When Fabric Manager logs into a Cisco MDS SAN-OS switch successfully through Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The SNMP v3 protocol data units (PDUs) with your Telnet/SSH login name as the SNMPv3 user are authenticated by the switch. Fabric Manager can temporarily use the Telnet/SSH login name as the SNMP v3 `auth` and `priv` passphrase. This temporary SNMP login is only allowed if you have one or more active MDS Shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMP v3 operations.

## Authorization

By default, two roles exist in all Cisco MDS switches:

- Network operator—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator—Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

If you use a SAN Volume Controller (SVC) setup, two more roles exist in all Cisco MDS switches:

- SVC administrator— Has permission to view the entire configuration and make SVC-specific configuration changes within the `switch(svc)` prompt.
- SVC operator—Has permission to view the entire configuration. The operator cannot make any configuration changes.



---

**Note** Refer to the *Cisco MDS 9000 Family SAN Volume Controller Configuration Guide* for more information on SVC.

---

These four default roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.



---

**Note** If a user only belongs to one of the newly-created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

---

## Accounting

The accounting feature tracks and maintains a log of every management session used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.



### Tip

---

The Cisco MDS 9000 Family switch uses interim-update RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have log update/watchdog packet flags in the AAA client configuration. Turn on this flag to ensure proper RADIUS accounting.

---



### Note

---

Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over AAA servers:

- It is easier to manage user password lists for each switch in the fabric.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- Easier to manage.
- Accounting log for all switches in the fabric can be centrally managed.
- Easier to manage user role mapping for each switch in the fabric.

## Remote Authentication Guidelines

When you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch. This is the recommended method.
- SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

## Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for fail-over servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fails to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

## AAA Service Configuration Options

AAA configuration in Cisco MDS 9000 Family switches is service based. You can have separate AAA configurations for the following services

- Telnet or SSH login (Cisco MDS Fabric Manager and Device Manager login).
- Console login.
- iSCSI authentication (see the [“iSCSI User Authentication”](#) section on page 20-17).
- FC-SP authentication (see [“Fibre Channel Security Protocol”](#) section on page 29-1).
- Accounting.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the methods fail, local is tried.

**Note**

Even if local is not specified as one of the options, it is tried when all other configured options fail.

## Configuring RADIUS

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

**Note**

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is click, the other Information pane tabs that use CFS are activated.

## Setting the RADIUS Server for Authentication and Accounting

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys.

By default, a switch retries a RADIUS server only once. This number can be configured. The maximum is five retries per server.

To add a RADIUS server, follow these steps:

- Step 1** Choose **Switches > Security > AAA** in Fabric Manager or choose **Security > AAA** in Device Manager.
- Step 2** Choose the **Servers** tab. You see the RADIUS or TACACS+ servers configured.
- Step 3** Click **Create Row** in Fabric Manager or **Create** in Device Manager. You see the Create Server dialog box.
- Step 4** Select the **radius** radio button to add a RADIUS server.
- Step 5** Set the IP address, authentication port and accounting port values.
- Step 6** Select whether the shared key is plain or encrypted in the KeyType field and set the key in the Key field.
- Step 7** Set the timeout and retry values for authentication attempts.
- Step 8** Click **Create** to create this RADIUS server or click **Close** to exit the dialog box without creating the new server.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Setting the Global Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when you create a new server.

To set the global preshared key, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Switches &gt; Security &gt; AAA</b> in Fabric Manager or choose <b>Security &gt; AAA</b> in Device Manager.  |
| <b>Step 2</b> | Choose the <b>Defaults</b> tab. You see the RADIUS and TACACS+ default settings.   |
| <b>Step 3</b> | Select whether the shared key is plain or encrypted in the KeyType field and set the key in the Key field.   |
| <b>Step 4</b> | Set the timeout and retry values for authentication attempts.  |
| <b>Step 5</b> | Click <b>Apply Changes</b> in Fabric Manager or <b>Apply</b> in Device Manager to save the global preshared key or click <b>Close</b> discard any unsaved changes. |
- 

## Defining Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named `cisco-avpair`. The value is a string with the following format:

```
protocol : attribute sep value *
```

Where `protocol` is a Cisco attribute for a particular type of authorization, and `sep` is = for mandatory attributes, and \* is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, like authorization information, along with authentication results. This authorization information is specified through VSAs.

### VSA Format

The following VSA protocol options are supported:

- `Shell` protocol—used in access-accept packets to provide user profile information.
- `Accounting` protocol—used in accounting-request packets. If a value contains any white spaces, it should be put within double quotation marks.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The following attributes are supported:

- **roles**—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles `vsan-admin` and `storage-admin`, the value field would be “`vsan-admin storage-admin`.” This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

```
shell:roles="network-admin vsan-admin"
shell:roles*"network-admin vsan-admin"
```

When an VSA is specified as `shell:roles*"network-admin vsan-admin"`, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- **accountinginfo**—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

## Specifying SNMPv3 on AAA Servers

The vendor/custom attribute `cisco-av-pair` can be used to specify user’s role mapping using the format:

```
shell:roles="roleA roleB ..."
```

As of Cisco MDS SAN-OS Release 2.0, the VSA format is enhanced to optionally specify your SNMPv3 authentication and privacy protocol attributes as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the `cisco-av-pair` attribute on the ACS server, MD5 and DES are used by default.



### Note

---

Only administrators can view the RADIUS preshared key.

---

## Configuring TACACS+

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## About TACACS+

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The addition of TACACS+ support in Cisco MDS SAN-OS Release 1.3(x) enables the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- TCP transport protocol to send data between the AAA client and server, using reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

## Enabling TACACS+

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. You must explicitly enable the TACACS+ feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

## Setting the TACACS+ Server

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server (see the [“Setting the Global Preshared Key”](#) section on page 26-6).

To add a TACACS+ server, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; Security &gt; AAA</b> in Fabric Manager or choose <b>Security &gt; AAA</b> in Device Manager.           |
| <b>Step 2</b> | Choose the <b>Servers</b> tab. You see the RADIUS or TACACS+ servers configured.  |
| <b>Step 3</b> | Click <b>Create Row</b> in Fabric Manager or <b>Create</b> in Device Manager. You see the Create Server dialog box.             |
| <b>Step 4</b> | Select the <b>tacacs+</b> radio button to add a RADIUS server.  |
| <b>Step 5</b> | Set the IP address, authentication port and accounting port values.   |
| <b>Step 6</b> | Select whether the shared key is plain or encrypted in the KeyType field and set the key in the Key field.                      |
| <b>Step 7</b> | Set the timeout and retry values for authentication attempts.   |
| <b>Step 8</b> | Click <b>Create</b> to create this TACACS+ server or click <b>Close</b> to exit the dialog box without creating the new server. |
- 

## Defining Custom Attributes for Roles

Cisco MDS 9000 Family switches use the TACACS+ custom attribute for service shells to configure roles to which a user belongs. TACACS+ attributes are specified in `name=value` format. The attribute name for this custom attribute is `cisco-av-pair`. The following example illustrates how to specify roles using this attribute:

```
cisco-av-pair=shell:roles="network-admin vsan-admin"
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

You can also configure optional custom attributes to avoid conflicts with non-MDS Cisco switches using the same AAA servers.

```
cisco-av-pair*shell:roles="network-admin vsan-admin"
```

Additional custom attribute shell:roles are also supported:

```
shell:roles="network-admin vsan-admin"
```

or

```
shell:roles*"network-admin vsan-admin"
```



**Note**

TACACS+ custom attributes can be defined on an Access Control Server (ACS) for various services (for example, shell). Cisco MDS 9000 Family switches require the TACACS+ custom attribute for the service shell to be used for defining roles.

## Supported TACACS+ Servers

The Cisco MDS SAN-OS software currently supports the following parameters for the listed TACACS+ servers:

- TACACS:

```
cisco-av-pair=shell:roles="network-admin"
```

- Cisco ACS TACACS

```
shell:roles="network-admin"
shell:roles*"network-admin"
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair*shell:roles*"network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

- Open TACACS

```
cisco-av-pair*shell:roles="network-admin"
cisco-av-pair=shell:roles*"network-admin"
```

## Configuring Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol: either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

You can configure these server groups at any time but they only take effect when you apply them to an AAA service.

To configure a RADIUS or TACACS+ server group, follow these steps:

- Step 1** Choose **Switches > Security > AAA** in Fabric Manager or choose **Security > AAA** in Device Manager.
- Step 2** Choose the **Server Group** tab. You see the RADIUS or TACACS+ servers configured.
- Step 3** Click **Create Row** in Fabric Manager or **Create** in Device Manager. You see the Create Server dialog box.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Select the **radius** radio button to add a RADIUS server group or select **tacacs+** to add a TACACS+ server group.
- Step 5** Check the servers from the ServerIdList for the servers you want to be part of this server group.
- Step 6** Click **Create** to create this RADIUS server or click **Close** to exit the dialog box without creating the new server.
- 

## Distributing AAA server Configuration

Configuration for RADIUS and TACACS+ AAA on a switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default.

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered there after are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.



### Note

Server group configurations are not distributed.

---

## Enabling the distribution

Only switches where distribution is enabled can participate in the distribution activity.

To enable a RADIUS or TACACS+ CFS distribution using Fabric Manager, follow these steps:

- Step 1** Choose **Switches > Security > AAA > RADIUS** or choose **Switches > Security > AAA > TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.
- Step 2** Choose the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration.
- Step 3** Choose **enable** from the Enable > Admin drop-down list for all switches that you want to enable CFS on for RADIUS or TACACS+.
- Step 4** Click **Apply Changes** to distribute these changes through the fabric.
- 

## Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a AAA configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

After you issue the first configuration command related to AAA servers, all server and global configuration made (including the configuration that caused the distribution session start) are stored in a temporary buffer—not in the running configuration.

## Committing the Distribution

The RADIUS or TACACS global and/or server configuration stored in the temporary buffer can be applied to the running configuration across all switches in the fabric (including the originating switch).

To distribute a RADIUS or TACACS+ configuration using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > AAA > RADIUS** or choose **Switches > Security > AAA > TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.
  - Step 2** Choose the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration
  - Step 3** Choose **commit** in the Config Changes > Action drop-down list for all switches that you want to enable CFS on for RADIUS or TACACS+.
  - Step 4** Click **Apply Changes** to distribute these changes through the fabric.
- 

## Discarding the Distribution Session

Discarding the distribution of a session-in-progress causes the configuration in the temporary buffer to be dropped. The distribution is not applied.

To discard a RADIUS or TACACS+ distribution using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > AAA > RADIUS** or choose **Switches > Security > AAA > TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.
  - Step 2** Choose the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration
  - Step 3** Choose **clear** from the Config Changes > Action drop-down list for all switches that you want to enable CFS on for RADIUS or TACACS+.
  - Step 4** Click **Apply Changes** to cancel the distribution.
- 

## Local AAA Services

The system maintains the user name and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information.

## Disabling AAA Authentication

You can turn off password verification. If you configure this option, users will be able to log in without giving a valid password. But the user should at least exist locally on the Cisco MDS 9000 Family switch.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Use this option cautiously. If configured, any user can to access the switch at any time.



## IP Access Control Lists

---

Cisco MDS 9000 Family IP storage services modules extend the reach of Fibre Channel SANs by using open-standard, IP-based technology. IP access control lists (IP-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IP-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

This chapter contains the following sections:

- [IP-ACL Configuration Guidelines, page 27-1](#)
- [Filter Contents, page 27-2](#)
- [Using the IP-ACL Wizard, page 27-4](#)
- [Creating Complex IP-ACLs Using Device Manager, page 27-5](#)
- [Associating IP-ACL Profiles to Interfaces, page 27-6](#)
- [Removing Associations Between IP-ACL Profiles and Interfaces, page 27-6](#)
- [Deleting IP Profiles, page 27-7](#)

## IP-ACL Configuration Guidelines

Each switch in the Cisco MDS 9000 Family can have a maximum of 64 IP-ACLs, and each IP-ACL can have a maximum of 256 filters. IP-ACLs can be associated with the management interface or any Gigabit Ethernet interface on the IP services modules (IPS-4, IPS-8, and MPS-14/2).

Follow these guidelines when configuring IP-ACLs in any switch or director in the Cisco MDS 9000 Family:

- In Cisco MDS SAN-OS Release 1.3 and earlier, you could only apply IP-ACLs to VSAN interfaces and the management interface. As of Cisco MDS SAN-OS Release 2.0(1b), you can also apply IP-ACLs to Gigabit Ethernet interfaces (IP services modules, including MPS-14/2 modules) and Ethernet PortChannel interfaces.



### Tip

---

If IP-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group.

---

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Caution**

Do not apply IP-ACLs to only one member of a PortChannel group. Apply IP-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IP-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.

## Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (TOS).

### Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, IP, Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

**Note**

When configuring IP-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

### Address Information

The address information is required in each filter. It identifies the following details:

- Source—The address of the network or host from which the packet is being sent.
- Source-wildcard—The wildcard bits applied to the source.
- Destination—The number of the network or host to which the packet is being sent.
- Destination-wildcard—The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IP address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IP address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 to require an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.
- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Port Information

The port information is optional. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. [Table 27-1](#) displays the port numbers recognized by the Cisco MDS SAN-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

**Table 27-1 TCP and UDP Port Numbers**

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514
TCP	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	wbem-http	5988
	wbem-https	5989


***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Using the IP-ACL Wizard

To use the IP-ACL Wizard to create an ordered list of IP filters in a named IP-ACL profile, follow these steps:

- Step 1** In Fabric Manager, choose the **IP-ACL Wizard** icon from the Fabric Manager toolbar. You see the IP-ACL Wizard.

**Figure 27-1** IP-ACL Wizard



- Step 2** Enter a **Name** for the IP-ACL profile.
- Step 3** Click the **Add** button to add a new rule to this IP-ACL profile. You see a new rule in the table with default values.
- Step 4** Modify the **Source Ip** and **Source Mask** as necessary for your filter.



**Note** The IP-ACL Wizard only creates inbound IP filters.

- Step 5** Choose the appropriate filter type from the Application column.
- Step 6** Choose **permit** or **deny** from the Action column.
- Step 7** Repeat [Step 3](#) through [Step 6](#) for additional IP filters.
- Step 8** Click **Up** or **Down** to order the filters in this IP-ACL profile.



**Tip** Order the IP filters carefully. Traffic is compared to the IP filters in order. The first match is applied and the rest are ignored.

- Step 9** Click **Next**. You see a list of switches that this IP-ACL profile can be applied to.
- Step 10** Uncheck any switches that you do not want this IP-ACL profile applied to.
- Step 11** Select the **Interface** you want this IP-ACL applied to.
- Step 12** Click **Finish** to create this IP-ACL profile and apply it to the selected switches, or click **Cancel** to exit the IP-ACL Wizard without creating an IP-ACL profile.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Creating Complex IP-ACLs Using Device Manager

The IP-ACL Wizard in Fabric Manager provides tools to create an ordered list of simple IP filters and apply those filters to switches in the fabric.

To create more complex IP-ACLs using Device Manager, follow these steps:

- 
- Step 1** Choose **Security > IP ACLs**. You see the IP-ACL dialog box.
  - Step 2** Click **Create ...** to create an IP-ACL profile.
  - Step 3** Enter a profile name and click **Create**. This creates an empty, named IP-ACL profile.
  - Step 4** Click on the IP-ACL profile you created and click **Rules...**. You see the list of IP filters associated with this profile.
  - Step 5** Click **Create...** to create an IP filter. You see the Create IP Filter dialog box.
  - Step 6** Choose the **permit** or **deny** Action radio button and set the Internet Protocol Number in the Protocol field. The drop-down menu provides common filtered protocols.
  - Step 7** Set the source IP address you want this filter to match against and the wildcard mask, or check the **Any** check box to match this filter against any IP address. This creates an IP filter that will check the source IP address of frames.



**Note** The wildcard mask denotes a subset of the IP Address you want to match against. This allows a range of addresses to match against this filter.

---

- Step 8** Set the transport layer source port range if the protocol chosen is TCP or UDP.
  - Step 9** Repeat [Step 7](#) and [Step 8](#) for the destination IP address and port range. This creates an IP filter that will check the destination IP address of frames.
  - Step 10** Set ToS, ICMPType, and ICMPCode as appropriate.
  - Step 11** Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.
  - Step 12** Check the **LogEnabled** check box if you want to log all frames that match this IP filter.
  - Step 13** Click **Create** to create this IP filter and add it to your IP-ACL profile or click **Close** to close the IP Filter dialog box without creating an IP filter.
- 

Any existing IP filters for this IP-ACL profile can be modified from the IP-ACL profiles dialog box but the filters cannot be reordered.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Associating IP-ACL Profiles to Interfaces

To associate the IP-ACL profile to an interface, follow these steps.

- 
- Step 1** From Fabric Manager, choose **Switches > Security > IP ACL** from the Physical Attributes pane. You see the IP-ACL configuration in the Information pane.
- From Device Manager, choose **Security > IP ACL**. You see the IP-ACL profiles dialog box.
- Step 2** Click the **Interfaces** tab.
- You see a list of interfaces and associated IP-ACL profiles.
- Step 3** Click the **Create Row** icon in Fabric Manager or the **Create** button in Device Manager. You see the Create Interface dialog box.
- Step 4** Optionally, select the switches you want to include in the IP-ACL profile by checking the check boxes next to the switch address in Fabric Manager.
- Step 5** Set the interface you want associated with an IP-ACL profile in the Interface field.
- Step 6** Choose the appropriate ProfileDirection radio button (either **inbound** or **outbound**).
- Step 7** Enter the profile name in the Profile Name field.



---

**Note** This profile name must already have been created using the Create Profiles dialog box. If not, no filters will be enabled until you go to the Create Profiles dialog box and create the profile.

---

- Step 8** Click **Create** to associate the profile, or click **Close** to close the Create Interfaces dialog box without associating a profile.
- You see the newly associated profile in the list of profiles.
- Step 9** Repeat [Step 8](#) to create additional associations, or click the **Close** button to close the Create Interfaces dialog box.
- 

## Removing Associations Between IP-ACL Profiles and Interfaces

To delete an IP-ACL profile, you must first delete all associations between that profile and the interfaces.

To remove associations between IP profiles and interfaces using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > IP ACL** from the Physical Attributes pane.
- You see the IP-ACL configuration in the Information pane.
- Step 2** Click the **Interfaces** tab.
- You see a list of switches, ACLs, and profile names.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 3** Select the row you want to delete. If you want to delete multiple rows, hold down the **Shift** key while selecting rows.
- Step 4** Click the **Delete Row** icon.
- The interfaces are disassociated from the profile.
- 

## Deleting IP Profiles

You must delete the association between IP profiles and interfaces before deleting the IP profile. To delete an IP profile using Fabric Manager, follow these steps.

- 
- Step 1** Choose **Switches > Security > IP ACL** from the Physical Attributes pane.
- You see the IP-ACL configuration in the Information pane.
- Step 2** Click the **Profiles** tab.
- You see a list of switches, ACLs, and profile names.
- Step 3** Select the row you want to delete. If you want to delete multiple rows, hold down the Shift key while selecting rows.
- Step 4** Click the **Delete Row** icon.
- The profiles are deleted.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## IPsec and IKE

---

Fabric Manager provides the capability to configure and manage IPsec using IKE.

This chapter includes the following sections:

- [Configuring IPsec Network Security, page 28-1](#)
- [Enabling IPsec Using FCIP Wizard, page 28-7](#)
- [Modifying IKE and IPsec, page 28-8](#)

## Configuring IPsec Network Security

IP Security Protocol (IPsec) is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides these security services at the IP layer. IPsec can be used to protect one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is per the latest version of RFC2401. Cisco SAN-OS IPsec implements RFC 2402 through RFC 2410.

Refer to the following website for further information on the IPsec RFCs:

<http://www.ietf.org/html.charters/ipsec-charter.html>.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys to be used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and additionally, implements the draft-ietf-ipsec-ikev2-15.txt draft.

Refer to the following website for further information on the IKE draft:

<http://www.ietf.org/internet-drafts/draft-ietf-ipsec-ikev2-15.txt>.



### Note

The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is also sometimes used to describe only the data services.

## The 14/2-Port Multiprotocol Services Module

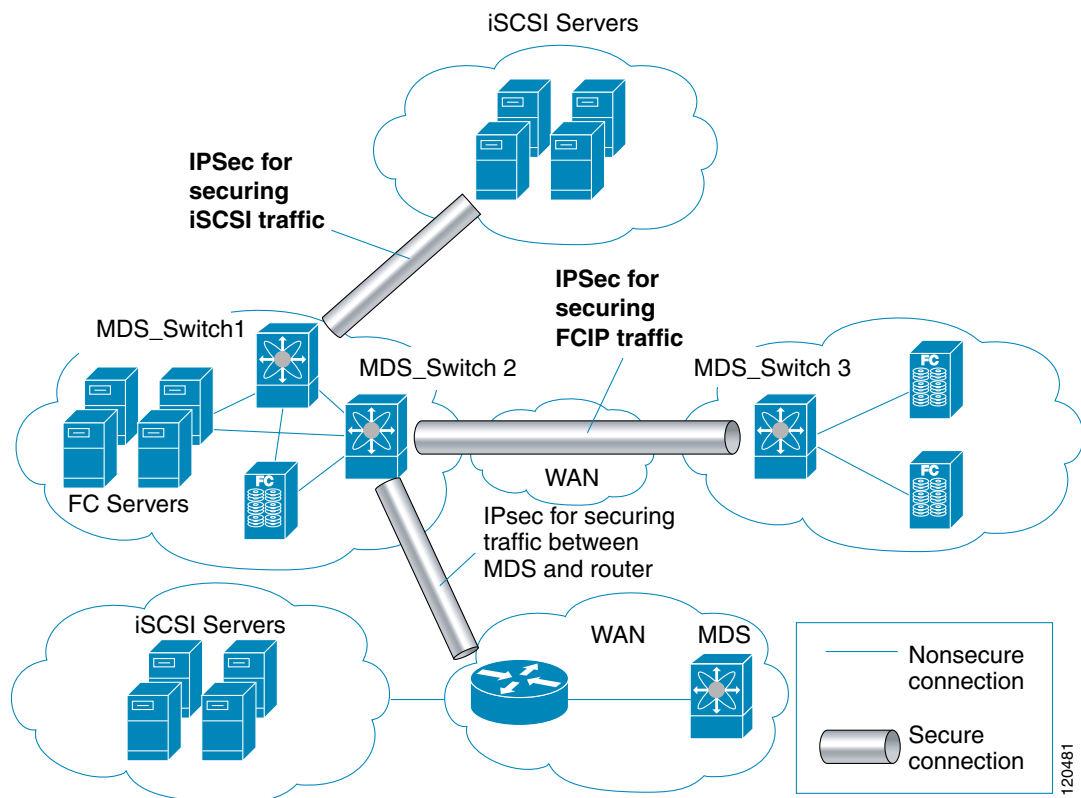
The 14/2-port Multiprotocol Services (MPS-14/2) module allows you to use Fibre Channel, FCIP, and iSCSI features. It integrates seamlessly into the Cisco MDS 9000 Family, and it supports the full range of features available on other switching modules, including VSANs, security, and traffic management.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

This module is available for use in any switch in the Cisco MDS 9200 Series or in the Cisco MDS 9500 Series. The 16-port, hot-swappable MPS-14/2 module has 14 Fibre Channel ports (numbered 1 through 14) and two Gigabit Ethernet ports (numbered 1 and 2) that can support FCIP protocol, iSCSI protocol, or both protocols simultaneously. The MPS-14/2 supports IPsec on the Gigabit Ethernet ports. See the “Enabling IPsec Using FCIP Wizard” section on page 28-7.

Figure 28-1 shows how the MPS-14/2 module is used in different scenarios.

**Figure 28-1 FCIP and iSCSI Scenarios Using MPS-14-2 Modules**



## IPsec Prerequisites

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE\_PKG license.
- Configure IKE.



### Note

The IPsec feature inserts new headers in existing packets.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## IPsec Compatibility

IPsec features are compatible with the following Cisco MDS hardware running Cisco MDS SAN-OS Release 2.0 or later:

- MPS-14/2 modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-Port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.



### Note

In both the MPS module and the Cisco MDS 9216i integrated supervisor module, the port numbering differs for the Fibre Channel and the Gigabit Ethernet ports—the Fibre Channel ports are numbered from 1 through 14 and the Gigabit Ethernet ports are numbered as 1 and 2.

IPsec features are compatible with the following fabric set up:

- Two connected Cisco MDS 9200 switches or Cisco MDS 9500 directors running Cisco MDS SAN-OS Release 2.0 or later.
- Cisco MDS 9200 switches or Cisco MDS 9500 directors running Cisco MDS SAN-OS Release 2.0 or later connected to any Cisco router.
- Cisco MDS 9200 switches or Cisco MDS 9500 directors running Cisco MDS SAN-OS Release 2.0 or later connected to any Cisco host.
- The following features are not supported in the SAN-OS implementation of the IPsec feature:
  - Authentication header (AH).
  - Transport mode.
  - Security association bundling.
  - Manually configuring security associations.
  - Per host security association option in a crypto map.
  - Security association idle timeout
  - Dynamic crypto maps.



### Note

Any reference to crypto maps in this document, only refers to static crypto maps.

## About IPsec

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec switches:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.

**Note**

The term data authentication is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco SAN-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.

## About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

Two versions of IKE are used in the SAN-OS implementation: IKE version 1 (IKEv1) and IKE version 2.

## IPsec and IKE Terminology

The terms used in this chapter are explained in this section.

- Security association (SA)—An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
  - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
  - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
  - Transform—A list of operations done on a dataflow to provide data authentication, data confidentiality, and data compression. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
  - Session keys—A key to encrypt and decrypt IP packets in a specified IKE session.
  - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and is automatically renegotiated (rekeyed).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Mode of operation—Two modes of operation are generally available for IPsec and IKE: tunnel mode and transport mode. The SAN-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet and an additional IP header between two hosts, a host and a gateway, or between two gateways. The gateways encrypt traffic on behalf of the hosts and subnets. This mode implements secure internal, external, remote access, and other networks. The SAN-OS implementation of IPsec does not support transport mode.



**Note** The term *tunnel mode* is different from the term *tunnel* used to indicate secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets in order to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.
- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
  - Data integrity—Verifies that data has not been altered.
  - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address/mask, destination address/mask, IP next protocol field, and source and destination ports, where the protocol and port fields can have the values of any. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.
- Perfect forwarding secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—an ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if should be allowed in clear text, or if it should be dropped.
  - IPsec SPDs are derived from user configuration of crypto maps.
  - IKE SPDs are configured by the user.

## Supported IPsec Transforms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode. This is an encryption technology.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet. This is an encryption technology.
- Triple DES (3DES) is a strong form of encryption that allows sensitive information to be transmitted over untrusted networks. It enables customers to utilize network layer encryption and implements 168-bit encryption. This is an encryption technology.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to [export@cisco.com](mailto:export@cisco.com).

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data. This is an authentication technology.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant. This is an authentication technology.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm. This is an authentication technology.

## Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol which allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) groups are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode. This is an encryption technology.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet. This is an encryption technology.
- Triple DES (3DES) is a strong form of encryption that allows sensitive information to be transmitted over untrusted networks. It enables customers to utilize network layer encryption and implements 168-bit encryption. This is an encryption technology.

**Note**

Cisco SAN-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to [export@cisco.com](mailto:export@cisco.com).

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data. This is an authentication technology.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant. This is an authentication technology.
- The switch authentication algorithm uses the preshared keys based on the IP address.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Supported Algorithms for Windows and Linux Platforms

Table 28-1 lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms.

**Table 28-1**      ***Supported Algorithms for Windows and Linux Platforms***

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5


## Enabling IPsec Using FCIP Wizard

Fabric Manager simplifies the configuration of IPsec and IKE by enabling and configuring these features as part of the FCIP configuration using the FCIP Wizard. See the “[Using the FCIP Wizard](#)” section on page 19-5.

To enable IPsec using Fabric Manager, follow these steps:

- Step 1**      Open the FCIP Wizard by clicking its icon in the Fabric Manager toolbar. [Figure 28-2](#) shows the FCIP Wizard icon.

**Figure 28-2**      ***FCIP Wizard***



- Step 2**      Choose the switches that act as endpoints for the FCIP link and click **Next**.

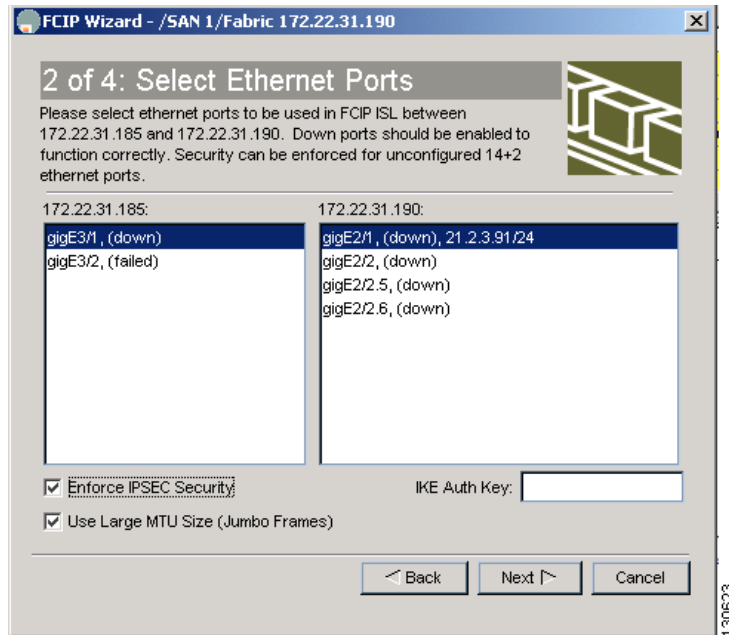


**Note**      These switches must have MPS-14/2 modules installed to configure IPsec on this FCIP link.

- Step 3**      Choose the Gigabit Ethernet ports on each MPS-14/2 module that will form the FCIP link.
- Step 4**      Check the **Enforce IPSEC Security** check box and set Ike Auth Key as shown in [Figure 28-3](#).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 28-3 Enabling IPsec on an FCIP Link**



- Step 5** Click **Next**. You see the TCP connection characteristics.
- Step 6** Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link. You can measure the round-trip time between the Gigabit Ethernet endpoints by clicking the **Measure** button.
- Step 7** Check the **Enable Write Acceleration** check box to enable FCIP write acceleration on this FCIP link. See the “[FCIP Write Acceleration](#)” section on page 19-4.
- Step 8** Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link. See the “[FCIP Compression](#)” section on page 19-5.
- Step 9** Click **Next** to configure the FCIP tunnel parameters.
- Step 10** Set the Port VSAN and click the **Trunk Mode** radio button for this FCIP link. See the “[Checking Trunk Status](#)” section on page 19-10.
- Step 11** Click **Finish** to create this FCIP link or click **Cancel** to exit the FCIP Wizard without creating an FCIP link.

## Modifying IKE and IPsec

Once IPsec is configured on an FCIP link, you can modify IKE and IPsec features using Fabric Manager. IKE must first be enabled and configured so the IPsec feature can trigger an SA with the required peer.

You cannot disable IKE if IPsec is enabled. When you disable the IKE feature, the IPsec configuration is cleared from the running configuration.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To verify that IPsec and IKE are enabled using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; Security &gt; IPSEC</b> in the Physical Attributes pane. You see the IPsec configuration in the Information pane. |
| <b>Step 2</b> | Choose the <b>Control</b> tab and verify that the switches you want to modify for IPsec are enabled in the Status column.                 |
| <b>Step 3</b> | Choose <b>Switches &gt; Security &gt; IKE</b> in the Physical Attributes pane. You see the IKE configuration in the Information pane.     |
| <b>Step 4</b> | Choose the <b>Control</b> tab and verify that the switches you want to modify for IKE are enabled in the Status column.                   |
- 

## Crypto ACL Guidelines

Follow these guidelines when configuring ACLs for the IPsec feature:

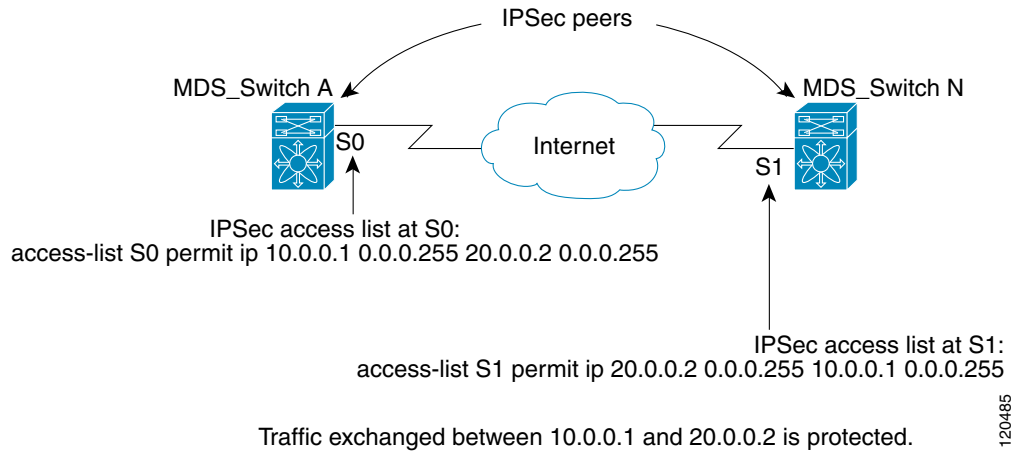
- The **permit** option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.
- The **deny** option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPsec ACL. Therefore, the ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.
- In [Figure 28-4](#), IPsec protection is applied to traffic between Host 10.0.0.1 and Host 20.0.0.2 as the data exits switch A's S0 interface enroute to Host 20.0.0.2. For traffic from Host 10.0.0.1 to Host 20.0.0.2, the ACL entry on switch A is evaluated as follows:
  - source = host 10.0.0.1
  - dest = host 20.0.0.2

For traffic from Host 20.0.0.2 to Host 10.0.0.1, that same ACL entry on switch A is evaluated as follows:

- source = host 20.0.0.2
- dest = host 10.0.0.1

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 28-4 IPsec Processing of Crypto ACLS**



- If you configure multiple statements for a given crypto ACL which is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.
- The IP ACLs used for traffic filtering purposes are also used for crypto.

## Mirror Image Crypto ACLs

For every crypto ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.



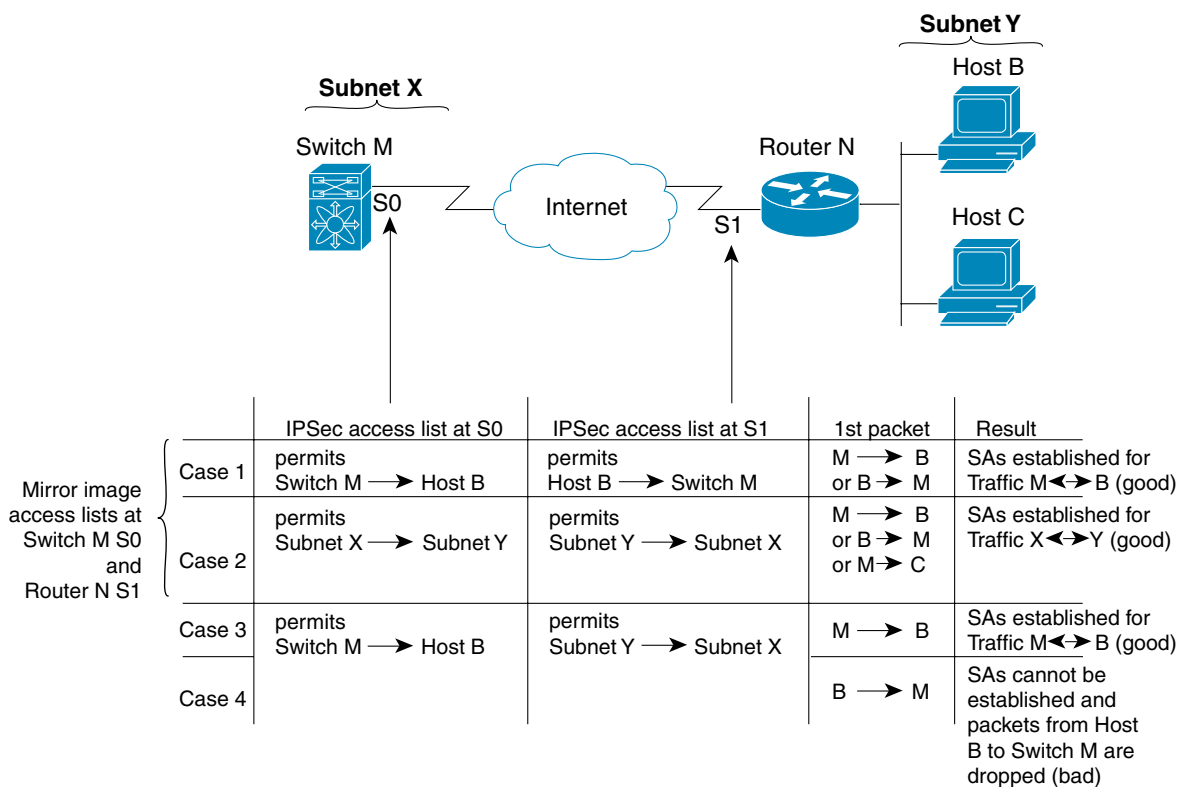
### Tip

The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.)

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

Figure 28-5 shows some sample scenarios with and without mirror image ACLs.

**Figure 28-5 IPsec Processing of Mirror Image Configuration**



As Figure 28-5 indicates, IPsec SAs can be established as expected whenever the two peers' crypto ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the ACLs are not mirror images of each other. This can happen in the case where an entry in one peer's ACL is a subset of an entry in the other peer's ACL, such as shown in Cases 3 and 4 of Figure 3. IPsec SA establishment is critical to IPsec—without SAs, IPsec does not work, causing any packets matching the crypto ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In Figure 28-5, an SA cannot be established in Case 4. This is because SAs are always requested according to the crypto ACLs at the initiating packet's end. In Case 4, switch N requests that all traffic between Subnet X and Subnet Y be protected, but this is a superset of the specific flows permitted by the crypto ACL at switch M so the request is therefore not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto ACL at switch N.

Because of the complexities introduced when crypto ACLs are not configured as mirror images at peer IPsec devices, Cisco strongly encourages you to use mirror image crypto ACLs.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## The any Keyword in Crypto ACLs



### Tip

We recommend that you configure mirror image crypto ACLs for use by IPsec and that you avoid using the **any** option.

The **any** option in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface—this configuration can cause multicast traffic to fail.

The **permit any any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use the **any** option in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

## Configuring Crypto IP-ACLs

You can configure IP-ACLs for crypto using the guidelines in the “[Crypto ACL Guidelines](#)” section on page 28-9.

See [Chapter 27, “IP Access Control Lists”](#) for guidelines on creating IP-ACLs using Fabric Manager.

## Transform Sets

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry’s access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers’ IPsec security associations.



### Tip

If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Table 28-2 provides a list of allowed transform combinations.

**Table 28-2 Allowed Transform Combinations**

Transform Type	Transform	Description
ESP encryption <sup>1</sup> transform (pick one.)	esp-des	ESP with the 56-bit DES encryption algorithm
	esp-3des	ESP with the 168-bit DES encryption algorithm (3DES or Triple DES)
	aes-128	In both counter mode <sup>2</sup> and CBC
	aes-256	
ESP authentication <sup>3</sup> transform (pick one.)	esp-md5-hmac	ESP with the MD5 (HMAC variant) authentication algorithm
	esp-sha-hmac	ESP with the SHA (HMAC variant) authentication algorithm
	aes-xcbc-mac	AES SCBC (MAC variant) ESP authentication algorithm

1. Mandatory.

2. If you select counter mode ESP encryption, authentication is required.

3. Optional in all other encryption cases (except counter mode).

## Crypto Map Entries

Once you have created the crypto ACLs, you can create crypto map sets to the interfaces. Crypto map IPsec entries pull together the various parts of the IPsec SA, including:

- The traffic to be protected by IPsec (per the crypto ACL). A crypto map set can contain multiple entries, each with a different ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets)
- Other parameters to define an IPsec SA

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A security policy database (SPD) is created for that interface
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decide whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto ACLs (for example, mirror image ACLs). If the responding peer entry is in the local crypto, the ACL must be permitted by the peer's crypto ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.
- If you create more than one crypto map entry for a given interface, use the `seq-num` of each map entry to rank the map entries: the lower the `seq-num`, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular ACL, the corresponding crypto map entry is tagged, and connections are established.

## The AutoPeer Option


Setting the peer address as AutoPeer in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be set up to each of the endpoints in the subnet specified by the crypto map's ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Figure 28-6 shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to set up SAs with the switch. Each host sets up its own SA, but shares the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Figure 28-6** *iSCSI with End-to-End IPsec Using the Auto-Peer Option*



120879

## SA Lifetime Negotiation

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.

## Perfect Forwarding Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forwarding secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

## Creating or Modifying Crypto Maps

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.
- Only one ACL is allowed for each crypto map entry (the ACL itself can have multiple entry or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the AutoPeer option to dynamically configure the peer.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To create or modify crypto map entries using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > IPSEC** in the Physical Attributes pane. You see the IPSEC configuration in the Information pane.
  - Step 2** Choose the **CryptoMap Set Entry** tab. You see the existing crypto maps configured.
  - Step 3** Optionally, click **Create Row** to create a new crypto map entry. You see the Create Crypto Map dialog box.
  - Step 4** Select the switch you want to configure or modify. If you are creating a new crypto map, set the setName and priority for this crypto map.
  - Step 5** Set the IP-ACL and TransformSetIdList for this crypto map.
  - Step 6** Optionally, check the **AutoPeer** check box or set the Peer address if you are creating a new crypto map. See the [“The AutoPeer Option” section on page 28-14](#).
  - Step 7** Choose the appropriate PFS radio button. See the [“Perfect Forwarding Secrecy” section on page 28-15](#).
  - Step 8** Set the Lifetime and LifeSize. See the [“SA Lifetime Negotiation” section on page 28-15](#).
  - Step 9** Optionally, click **Create** if you are creating a new crypto map, or click the **Apply Changes** icon if you are modifying an existing crypto map.
- 

## Applying a Crypto Map Set to an Interface

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

To apply a crypto map set to an interface using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > IPSEC** in the Physical Attributes pane. You see the IPSEC configuration in the Information pane.
  - Step 2** Choose the **Interfaces** tab. You see the existing interface to crypto map configuration.
  - Step 3** Optionally, click **Create Row** to create a apply a crypto map to an interface. You see the Interfaces Create dialog box.
  - Step 4** Select the switch and interface you want to configure.
  - Step 5** Select the **CryptomapSetName** to the name of the crypto map you want to apply to this interface.
  - Step 6** Click **Create** to apply the crypto map to the selected interface or click **Close** to exit the dialog box without applying the crypto map.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be re-established with the changed configuration. If the switch is actively processing IPsec traffic, it is desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.

## Global Lifetime Values

You can change the global lifetime values which are used when negotiating new IPsec SAs and override configured global lifetime values for a specified crypto map entry.

You can configure two lifetimes: timed or traffic-volume. A SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 4,500 MB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKE version 1 (IKEv1) to setup IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKE version 2 (IKEv2) to setup IPsec SAs, SAs on each end has its own set up of lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold (when 10% of the configured value still remains) of the existing SA is reached, to ensure that negotiation completes before the existing SA expires.

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## FC-SP and DHCHAP

---

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. It consists of the CHAP protocol combined with the Diffie-Hellman exchange.

This chapter includes the following sections:


- [Fibre Channel Security Protocol, page 29-1](#)
- [Configuring DHCHAP Authentication, page 29-3](#)

## Fibre Channel Security Protocol

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics. For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see [Figure 29-1](#)).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 29-1 Switch and Host Authentication**



105209



**Note**

Fibre Channel host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

## About DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, thus preventing unauthorized devices from accessing the switch.



**Note**

The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and Diffie-Hellman groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE\_PKG license (see [Chapter 9, “Obtaining and Installing Licenses”](#)).

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## DHCHAP Compatibility with Existing Cisco MDS Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

## Configuring DHCHAP Authentication

To configure DHCHAP authentication using the local password database, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Enable DHCHAP.   |
| <b>Step 2</b> | Identify and configure the DHCHAP authentication modes.                              |
| <b>Step 3</b> | Configure the hash algorithm and DH group.   |
| <b>Step 4</b> | Configure the DHCHAP password for the local switch and other switches in the fabric. |
| <b>Step 5</b> | Configure the DHCHAP timeout value for reauthentication.                             |
| <b>Step 6</b> | Verify the DHCHAP configuration.   |
- 

## Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

To enable DHCHAP and FC-SP, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | From Fabric Manager, choose <b>Switches &gt; Security &gt; FC-SP</b> . You see the FC-SP configuration in the Information pane.<br><br>From Device Manager, choose <b>Security &gt; FC-SP</b> . You see the FC-SP Enable dialog box. Click <b>Yes</b> to enable FC-SP and DHCHAP for this switch. |
| <b>Step 2</b> | Choose the <b>Control</b> tab in Fabric Manager. You see the FC-SP enable state for all switches in the fabric.   |
| <b>Step 3</b> | Set the Command drop-down menu to <b>enable</b> for all switches that you want to enable FC-SP on.  |

[Send documentation comments to mdsfeedback-doc@cisco.com.](mailto:mdsfeedback-doc@cisco.com)

**Step 4** Click the **Apply Changes** icon to enable FC-SP and DHCHAP on the selected switches.

## Configuring DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode. When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- **On**—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- **AutoActive**—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- **AutoPassive (default)**—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- **Off**—The switch does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.



**Note**

Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 29-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

**Table 29-1 DHCHAP Authentication Status Between Two MDS Switches**

Switch N DHCHAP Modes	Switch 1 DHCHAP Modes			
	on	auto-active	auto-passive	off
on	FC-SP authentication is performed.	FC-SP authentication is performed.	FC-SP authentication is performed.	Link is brought down.
auto-Active			FC-SP authentication is <i>not</i> performed.	FC-SP authentication is <i>not</i> performed.
auto-Passive				
off	Link is brought down.	FC-SP authentication is <i>not</i> performed.		



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To configure the DHCHAP port authentication mode, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Interfaces > FC Physical**. You see the FC-SP configuration in the Information pane.
- From Device Manager, choose **Security > FC-SP**. You see the FC-SP Configuration dialog box.
- Step 2** Choose the **FC-SP** tab. You see the DHCHAP authentication mode for each interface.
- Step 3** Set the Mode drop-down menu to the DHCHAP authentication mode you want to configure for that interface.
- Step 4** Click the **Apply Changes** icon in Fabric Manager or click **Apply** in Device Manager to save these DHCHAP port mode settings.
- 

## Changing the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.



### Tip

If you change the hash algorithm priority list, then change it globally for all switches in the fabric.



### Caution

RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

To change the DHCHAP hash algorithm priority list using Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Security > FC-SP** in Fabric Manager. You see the FC-SP configuration in the Information pane.
- Step 2** Choose the **General/Password** tab. You see the DHCHAP general settings mode for each switch.
- Step 3** Change the HashList for each switch in the fabric.
- Step 4** Click the **Apply Changes** icon to save the updated hash algorithm priority list or click the **Undo Changes** icon to discard any unsaved changes.
- 

## Changing DHCHAP Group Settings

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.



### Tip

If you change the DH group configuration, change it globally for all switches in the fabric.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To change the DH group list using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>Switches &gt; Security &gt; FC-SP</b> in Fabric Manager. You see the FC-SP configuration in the Information pane.                  |
| <b>Step 2</b> | Choose the <b>General/Password</b> tab. You see the DHCHAP general settings mode for each switch.  |
| <b>Step 3</b> | Change the GroupList for each switch in the fabric.  |
| <b>Step 4</b> | Click the <b>Apply Changes</b> icon to save the updated DH group lists or click the <b>Undo Changes</b> icon to discard any unsaved changes. |
- 

## Configuring the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric.
- Approach 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Approach 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.



### Note

All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.



### Tip

We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family Fabric Manager to manage the password database.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring the DHCHAP Password for the Local Switch

To configure the DHCHAP password for the local switch, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Security > FC-SP**. You see the FC-SP configuration in the Information pane.  
From Device Manager, choose **Security > FC-SP**. You see the FC-SP Enable dialog box.
  - Step 2** Choose the **Local Passwords** tab. You see the DHCHAP local password for each switch.
  - Step 3** Click the **Create Row** icon in Fabric Manager or **Create** in Device Manager to create a new local password. You see the Create Local Passwords dialog box.
  - Step 4** Optionally, check the switches that you want to configure the same local password on in Fabric Manager.
  - Step 5** Select the switch **WNN** and set the Password.
  - Step 6** Click the **Create** to save the updated password or click **Close** to discard any unsaved changes.
- 

## Configuring Remote Passwords for Other Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).



### Note

The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

To configure the DHCHAP password for remote switches, follow these steps:

- 
- Step 1** From Fabric Manager, choose **Switches > Security > FC-SP**. You see the FC-SP configuration in the Information pane.  
From Device Manager, choose **Security > FC-SP**. You see the FC-SP Enable dialog box.
  - Step 2** Choose the **Remote Passwords** tab. You see the DHCHAP local password for each switch.
  - Step 3** Click the **Create Row** icon in Fabric Manager or **Create** in Device Manager to create a remote password. You see the Create Remote Passwords dialog box.
  - Step 4** Optionally, check the switches that you want to configure the same remote password on in Fabric Manager.
  - Step 5** Select the switch **WNN** and set the Password.
  - Step 6** Click the **Create** to save the updated password or click **Close** to discard any unsaved changes.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Setting the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.
- The same value must also be configured all switches in the fabric.

To change the DHCHAP timeout value using Fabric Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose <b>Switches &gt; Security &gt; FC-SP</b> in Fabric Manager. You see the FC-SP configuration in the Information pane.                 |
| <b>Step 2</b> | Choose the <b>General/Password</b> tab. You see the DHCHAP general settings mode for each switch.   |
| <b>Step 3</b> | Change the DHCHAP timeout value for each switch in the fabric.  |
| <b>Step 4</b> | Click the <b>Apply Changes</b> icon to save the updated timeout value or click the <b>Undo Changes</b> icon to discard any unsaved changes. |
- 

## Configuring DHCHAP AAA Authentication

You can individually set authentication options. If authentication is not configured, local authentication is used by default.

## Enabling FC-SP on ISLs

There is a new ISL pop-up menu called Enable FC-SP that enables FC-SP on switches at either end of the ISL. You are prompted for an FC-SP generic password, then asked to set FC-SP interface mode to ON for affected ports. Right-click an ISL and click **Enable FC-SP** to access this feature.



## Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.



**Note**

Port Security is only supported for Fibre Channel ports.

This chapter includes the following sections:

- [About Port Security, page 30-1](#)
- [Configuring Port Security, page 30-3](#)
- [Port Security Manual Configuration, page 30-6](#)

## About Port Security

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.

## About Auto-Learn

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature to activate the port security feature for the first time as it saves manual configuration for each port. Auto-learn is configured on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Auto-Learning Device Authorization

Table 30-1 summarizes the authorized connection for device requests.

**Table 30-1 Auto-Learn Device Authorization**

Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization	Condition
Configured with one or more switch ports	A switch on configured ports	Permitted	1
	A switch on other ports	Denied	2
Not configured	A port that is not configured	Permitted if auto-learn enabled	3
		Denied if auto-learn disabled	4
Configured or not configured	A switch port that allows any device	Permitted	5
Configured to log in to any switch port	Any port on the switch	Permitted	6
Not configured	A port configured with some other device	Denied	7

## Port Security Enforcement

If you choose to manually configure port security, you must configure the devices and switch port interfaces through which each device or switch is connected:

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Configuring Port Security

To configure port security, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Enable the port security feature on all participating switches. See the “ <a href="#">Enabling Port Security</a> ” section on page 30-3.   |
| <b>Step 2</b> | Activate with auto-learning, if you want to use the auto-learning feature to populate the port security database. See the “ <a href="#">Activating Port Security with Auto-Learn</a> ” section on page 30-3. |
| <b>Step 3</b> | Optionally, manually configure the port security database and then activate it. See the “ <a href="#">Manually Configuring Port Security</a> ” section on page 30-7.   |
- 

## Enabling Port Security

Before you can activate port security, you need to enable this feature on all switches in the fabric that will participate in port security.

To enable port security using Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>VSANxxx &gt; Port Security</b> from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane. |
| <b>Step 2</b> | Click the <b>CFS</b> tab and enable CFS on all participating switches in the VSAN.   |
| <b>Step 3</b> | Click the <b>Apply Changes</b> icon to enable CFS distribution for the port security feature.  |
| <b>Step 4</b> | Click the <b>Control</b> tab. You see the port security enable state for all switches in the selected VSAN.  |
| <b>Step 5</b> | Set the command column to <b>enable</b> for each switch in the VSAN.   |
| <b>Step 6</b> | Click the <b>CFS</b> tab and set the command column to <b>commit</b> on all participating switches in the VSAN.  |
| <b>Step 7</b> | Click the <b>Apply Changes</b> icon to distribute the port security enable to all switches in the VSAN.  |
- 

## Activating Port Security with Auto-Learn

To activate port security with auto-learn, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | From Fabric Manager, choose <b>VSANxxx &gt; Port Security</b> from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.<br>From Device Manager, <b>Choose Security &gt; Port...</b> You see the Port Security dialog box.  |
| <b>Step 2</b> | Click the <b>Actions</b> tab.  |
| <b>Step 3</b> | Click the <b>Action</b> column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options: <ul style="list-style-type: none"><li>• activate—Valid port security settings are activated.</li><li>• activate (TurnLearningOff)—Valid port security settings are activated and autolearn turned off.</li><li>• forceActivate—Activation is forced.</li></ul> |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **forceActivate(TurnLearningOff)**—Activation is forced and autolearn is turned off.
- **deactivate**—All currently active port security settings are deactivated.
- **NoSelection**—No action is taken.

- Step 4** Select the option you want to specify a port security setting action for that switch.
- Step 5** Check the **AutoLearn** check box for each switch in the VSAN to enable auto-learning. Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.
- Step 6** Click the **Apply Changes** icon in Fabric Manager or **Apply** in Device Manager to save these changes or click **Undo Changes** in Fabric Manager or **Close** in Device Manager to discard any unsaved changes.
- 

## Displaying Activated Port Security Settings

To display active port security settings, follow these steps:

- Step 1** Choose **VSANxxx > Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Active Database** tab.
- You see the active port security settings for that VSAN.
- 

## Displaying Port Security Statistics

To display port security statistics, follow these steps:

- Step 1** Choose **VSANxxx > Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Statistics** tab. You see the port security statistics for that VSAN.
- 

## Displaying Port Security Violations

Port violations are invalid login attempts (for example, login requests from unauthorized Fibre Channel devices). You can display a list of these attempts on a per-VSAN basis, using Fabric Manager.

To display port security violations, follow these steps:

- Step 1** Choose **VSANxxx > Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Violations** tab. You see the port security violations for that VSAN.
-



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Turning Auto-Learning On or Off

To turn auto-learning on or off, follow these steps:

- 
- Step 1** Choose **VSANxxx > Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.
- Step 2** Click the **Action** tab. You see the switches for that VSAN.
- Step 3** Check the **AutoLearn** check box next to the switch if you want to enable auto-learning.
- Step 4** Uncheck the **AutoLearn** check box next to the switch if you want to disable auto-learning.
- Step 5** Click the **CFS** button at the top of the Information pane and select **commit** .
- Step 6** Click the **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.
- 

## Example of Port Security Authorization

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 30-2 summarizes the port security authorization results for this active database.

**Table 30-2 Authorization Results for Scenario**

Scenario	Device Connection Request	Authorization	Condition	Reason
1	P1, N2, F1	Permitted	1	No conflict.
2	P2, N2, F1	Permitted	1	No conflict.
3	P3, N2, F1	Denied	2	F1 is bound to P1/P2.
4	P1, N3, F1	Permitted	6	Wildcard match for N3.
5	P1, N1, F3	Permitted	5	Wildcard match for F3.
6	P1, N4, F5	Denied	2	P1 is bound to F1.
7	P5, N1, F5	Denied	2	N1 is only allowed on F2.
8	P3, N3, F4	Permitted	1	No conflict.
9	S1, F10	Permitted	1	No conflict.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 30-2 Authorization Results for Scenario (continued)**

Scenario	Device Connection Request	Authorization	Condition	Reason
10	S2, F11	Denied	7	P10 is bound to F11.
11	P4, N4, F5 (auto-learn on)	Permitted	3	No conflict.
12	P4, N4, F5(auto-learn off)	Denied	4	No match.
13	S3, F5 (auto-learn on)	Permitted	3	No conflict.
14	S3, F5 (auto-learn off)	Denied	4	No match.
15	P1, N1, F6 (auto-learn on)	Denied	2	P1 is bound to F1.
16	P5, N5, F1 (auto-learn on)	Denied	7	P3 is bound to F1.
17	S3, F4 (auto-learn on)	Denied	7	P3 paired with F4.
18	S1, F3 (auto-learn on)	Permitted	5	No conflict.
19	P5, N3, F3	Permitted	6	Wildcard match for F3 and N3.
20	P7, N3, F9	Permitted	6	Wildcard match for N3.

## Port Security Manual Configuration



### Note

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is clicked, the other tabs in the Information pane that use CFS are activated.

To manually configure port security on any switch in the Cisco MDS 9000 Family, follow these steps:

- Step 1** Identify the WWN of the ports that need to be secured.
- Step 2** Secure the fWWN to an authorized nWWN or pWWN.
- Step 3** Activate the port security database.
- Step 4** Verify your configuration.

## WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port:
  - is allowed to login to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
  - nWWN is bound to a Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.
- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

## Manually Configuring Port Security

To manually configure port security on a switch, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose VSANxxx > <b>Port Security</b> from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.                  |
| <b>Step 2</b> | Click the <b>Config Database</b> tab. You see the configured port security settings for that VSAN.   |
| <b>Step 3</b> | Click the <b>Create Row</b> icon to add an authorized port pair. You see the Create Port Security dialog box.  |
| <b>Step 4</b> | Double-click the device from the available list for which you want to create the port security setting.  |
| <b>Step 5</b> | Double-click the port from the available list to which you want to bind the device.  |
| <b>Step 6</b> | Click <b>Create</b> to creating the port security setting, or click <b>Close</b> to close the Create Port Setting dialog without adding a new port security setting. |
| <b>Step 7</b> | Click the <b>CFS</b> tab and set the command column to <b>commit</b> on all participating switches in the VSAN.  |
| <b>Step 8</b> | Click the <b>Apply Changes</b> icon to save these changes or click <b>Undo Changes</b> to discard any unsaved changes.   |
- 

## Deleting a Port Security Pair

To delete a port security setting from the configured database on a switch, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Choose VSANxxx > <b>Port Security</b> from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane. |
| <b>Step 2</b> | Click the <b>Config Database</b> tab. You see the configured port security settings for that VSAN.  |
| <b>Step 3</b> | Click the row you want to delete.   |
| <b>Step 4</b> | Click the <b>Delete Row</b> icon. You see the confirmation dialog box.  |
| <b>Step 5</b> | Click <b>Yes</b> to delete the row, or click <b>No</b> to close the confirmation dialog box without deleting the row.                               |
| <b>Step 6</b> | Click the <b>CFS</b> tab and set the command column to <b>commit</b> on all participating switches in the VSAN.                                     |
| <b>Step 7</b> | Click the <b>Apply Changes</b> icon to save these changes or click <b>Undo Changes</b> to discard any unsaved changes.                              |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Database Interaction

Table 30-3 lists the differences and interactions between the active and configuration databases.

**Table 30-3      *Active and Configuration Port Security Databases***

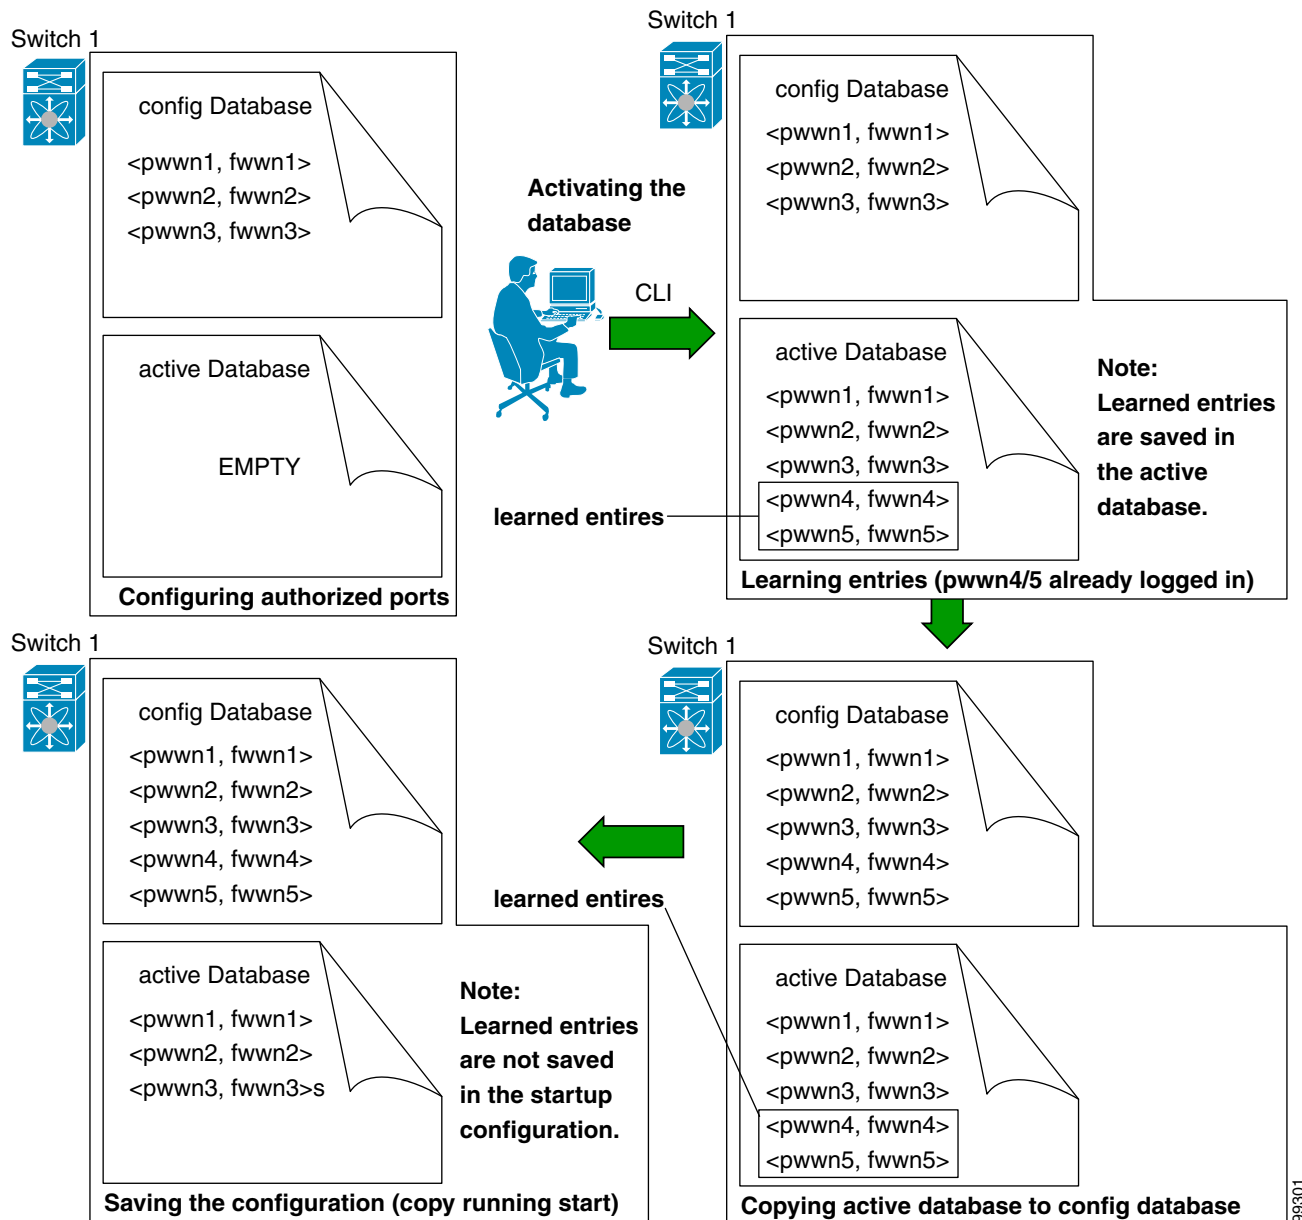
Configuration Database	Active Database
Read-write.	Read-only.
Saving the configuration saves all the entries in the configuration database.	Saving the configuration only saves the activated entries. Learned entries are not saved.
Once activated, the configuration database can be modified without any effect on the active database.	Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.
You can overwrite the configuration database with the active database.	You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

## Database Scenarios

The various scenarios in Figure 30-1 depict the active database and the configuration database status based on port security configurations.

**Figure 30-1 Port Security Database Scenarios**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Activating the Port Security Database

When you activate the port security database, all entries in the configured database are copied to the active database. After the database is activated, subsequent device login is subject to the activated port bound WWN pairs. Additionally, all devices that have already logged into the VSAN at the time of activation are also learned and added to the active database. If the auto-learn feature is already enabled in a VSAN, you will not be allowed to activate the database.

To activate port security with auto-learn disabled, follow these steps:

- 
- Step 1** From Fabric Manager, choose VSANxxx > **Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.
- From Device Manager, **Choose Security > Port...** You see the Port Security dialog box.
- Step 2** Click the **Actions** tab.
- Step 3** Click in the **Action** column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:
- activate—Valid port security settings are activated.
  - activate (TurnLearningOff)—Valid port security settings are activated and autolearn turned off.
  - forceActivate—Activation is forced.
  - forceActivate(TurnLearningOff)—Activation is forced and autolearn is turned off.
  - deactivate—All currently active port security settings are deactivated.
  - NoSelection— No action is taken.
- Step 4** Set the Action field you want for that switch.
- Step 5** Uncheck the **AutoLearn** check box for each switch in the VSAN to disable auto-learning.
- Step 6** Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.
- Step 7** Click the **Apply Changes** icon in Fabric Manager or **Apply** in Device Manager to save these changes or click **Undo Changes** in Fabric Manager or **Close** in Device Manager to discard any unsaved changes.
- 

## Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- If the auto-learn feature was enabled before the activation. To reactivate a database in this state.
- The exact security is not configured for each PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Forceful Port Security Activation

If the auto-learn option was enabled before the activation, reactivate the database. If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed with the activation by using the forceActivate option.

See the [“Activating the Port Security Database” section on page 30-10](#).



### Note

An activation using the forceActivate option logs out existing devices if they violate the active database.

## Database Reactivation



### Tip

If the auto-learn option is enabled and you activate the database, you will not be allowed to proceed.

To reactivate the database using Fabric Manager, follow these steps:

**Step 1** Disable auto-learning.

**Step 2** Copy the active database to the configured database.



### Tip

If the active database is empty, you cannot perform this step.

**Step 3** Activate the database.

## Copying an Active Database to the Config Database

To copy the active database to the config database using Fabric Manager, follow these steps:

**Step 1** Choose VSANxxx > **Port Security** from the Logical Domains pane. You see the port security configuration for that VSAN in the Information pane.

**Step 2** Click the **Actions** tab. You see the switches for that VSAN.

**Step 3** Check the **CopyActive ToConfig** check box next to the switch for which you want to copy the database. The active database is copied to the config database when the security setting is activated.

**Step 4** Uncheck the check box if you do not want the database copied when the security setting is activated.

**Step 5** Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

**Step 6** Click the **Apply Changes** icon to save these changes or click **Undo Changes** to discard any unsaved changes.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 5**

### **Network and Performance Monitoring**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## Network Monitoring

---

The primary purpose of Fabric Manager is to manage the network. In particular, SAN discovery and network monitoring are two of its key network management capabilities.

This chapter contains the following sections:

- [SAN Discovery and Topology Mapping, page 31-1](#)
- [Configuring System Message Logging, page 31-4](#)
- [Health and Event Monitoring, page 31-9](#)

## SAN Discovery and Topology Mapping

Fabric Manager provides extensive SAN discovery, topology mapping, and information viewing capabilities. Fabric Manager collects information on the fabric topology through SNMP queries to the switches connected to it. Fabric Manager recreates a fabric topology, presents it to the user in a customizable map, and provides inventory and configuration information in multiple viewing options.

### Device Discovery

Once Fabric Manager is invoked, a SAN discovery process begins. Using information polled from a seed Cisco MDS 9000 Family switch, including Name Server registrations, Fibre Channel Generic Services (FC-GS), Fabric Shortest Path First (FSPF), and SCSI-3, Fabric Manager automatically discovers all devices and interconnects on one or more fabrics. All available switches, host bus adapters (HBAs), and storage devices are discovered. The Cisco MDS 9000 Family switches use Fabric-Device Management Interface (FMDI) to retrieve HBA model, serial number and firmware version, and host operating-system type and version discovery without host agents. Fabric Manager gathers this information through SNMP queries to each switch. The device information discovered includes device names, software revision levels, vendor, ISLs, PortChannels, and VSANs.

### Topology Mapping

Fabric Manager is built upon a topology representation of the fabric. Fabric Manager provides an accurate view of multiple fabrics in a single window by displaying topology maps based on device discovery information. The user may modify the topology map icon layout with an easy-to-use, drag-and-drop interface. The topology map visualizes device interconnections, highlights configuration

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

information such as zones, VSANs, and ISLs exceeding utilization thresholds. The topology map also provides a visual context for launching command-line interface (CLI) sessions, configuring PortChannels, and opening device managers.

## Using the Topology Map

The Fabric Manager topology map can be customized to provide a view into the fabric that varies from showing all switches, end devices, and links, to showing only the core switches with single bold lines for any multiple links between switches. Use the icons along the left side of the topology map to control these views or right-click anywhere in the topology map to access the map controls.

You can zoom in or out on the topology map to see an overview of the SAN or focus on an area of importance. You can also open an overview window that shows the entire fabric. From this window, you can right-click and draw a box around the area you want to view in the main topology map view.

Another way to limit the scope of the topology display is to select a fabric or VSAN from the Logical Domains pane. The topology map displays only that fabric or VSAN.

Moving the mouse pointer over a link or switch provides a simple summary of that SAN component, along with a status indication. Right-clicking on the component brings up a pop-up menu. You can view the component in detail or access configuration or test features for that component.

Double-click a link to bring link status and configuration information to the information pane.

Double-click a switch to bring up Device Manager for that switch.

## Saving a Customized Topology Map Layout

Changes made to the topology map can be saved so that the customized view is available any time you open the Fabric Manager client for that fabric.

To save the customized layout in Fabric Manager, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Choose <b>File &gt; Preferences</b> to open the Fabric Manager preferences dialog box.                                     |
| <b>Step 2</b> | Click the <b>Map</b> tab and check the <b>Automatically Save Layout</b> check box to save any changes to the topology map. |
| <b>Step 3</b> | Click <b>Apply</b> to save this change, or click <b>Close</b> to discard any unsaved changes and close the dialog box.     |
- 

## Using Enclosures with Fabric Manager Topology Maps

Because not all devices are capable of responding to FC-GS-3 requests, different ports of a single server or storage subsystem may be displayed as individual end devices on the topology map. See the [“Modifying Device Grouping” section on page 3-14](#) to group these ports into a single enclosure for Fabric Manager.

The Alias->Enclosure button displays in the Information pane for hosts and storage elements. This button acts as a shortcut to naming enclosures. To use this shortcut, you highlight each row in the host or storage table that you want grouped in an enclosure and then click **Alias -> Enclosure**. This automatically sets the enclosure names of each selected row with the first token of the alias.


***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Mapping Multiple Fabrics

To log into multiple fabrics, the same username and password must be used. The information for both fabrics is displayed, with no need to select a seed switch. To see details of a fabric, click the tab for that fabric at the bottom of the Fabric pane, or double-click the fabric's cloud icon.

When you quit the Fabric Manager client, you can have Fabric Manager Server continuously monitor that fabric. Alternatively, you can use Fabric Manager client to select a fabric to monitor.

To continuously monitor a fabric in Fabric Manager, follow these steps:

- 
- Step 1** Choose **Server > Admin**. You see the Server Admin dialog box with a list of fabrics.
- Step 2** Check the **Continuously Monitor** check box next to the fabric(s) you want Fabric Manager Server to monitor.
- Step 3** Click **Apply**.
- The Continuously Monitor feature requires the purchase of the Fabric Manager Server license package. If you have not purchased and installed this package, you see a pop-up window informing you that you are about to enable a demo license for this feature. Click **Yes** to enable the demo license.
-  **Note** When you are finished checking out the demo, you can “check in” the feature by clicking the **Check In FM** button as described in the [“Fabric Manager Server Licensing” section on page 9-12](#).
- 
- Step 4** Click **Close** to close the Server Admin dialog box.
- 

## Inventory Management

The Information pane in Fabric Manager shows inventory, configuration, and status information for all switches, links, and hosts in the fabric. Inventory management includes vendor name and model, and software or firmware versions. Select a fabric or VSAN from the Logical Domains pane, and then select the **Summary** tab in the Information pane to get a count of the number of VSANS, switches, hosts, and storage elements in the fabric. See the [“Using Fabric Manager Client” section on page 3-3](#) for more information on the Fabric Manager user interface.

### Using the Inventory Tab from Fabric Manager Web Services

If you have configured Fabric Manager Web Services, you can launch this application and access the Inventory tab to see a summary of the fabrics managed by the Fabric Manager Server. The Inventory tab shows an inventory of the selected SAN, fabric, or switch.

**Summary**—Shows all VSANs, switches, and ports in the selected SAN or fabric.

**VSANs**— Shows all VSANs in the selected SAN or fabric.

**Switches**—Shows all attributes (such as IP address, vendor, and model) for all switches in the selected SAN, fabric, or VSAN.

**Modules**—Shows all line cards, fans, and power supplies for all switches in the selected SAN, fabric, or VSAN.

**ISLs**—Shows all the Inter-Switch Links for the selected SAN, fabric, or VSAN.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Hosts**—Shows all the host ports on the selected SAN, fabric, or VSAN.

**Storage**—Shows all the storage ports on the selected SAN, fabric, or VSAN.

**Zones**—Shows all the active zone members (including those in inter-VSAN zones) for the selected SAN, fabric, or VSAN.

See [Chapter 5, “Fabric Manager Web Services”](#) for more information on how to configure and use Fabric Manager Web Services.

## Configuring System Message Logging

The system message logging software saves messages in a log file or directs the messages to other devices. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting.
- Allows you to select the types of captured logging information.
- Allows you to select the destination of the captured logging information.

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. You can specify which system messages should be saved based on the type of facility and the severity level. See the “[Syslog Server Logging Facilities and Severity Levels](#)” section on [page 31-4](#). Messages are time-stamped to enhance real-time debugging and management.

The switch software saves system messages in a file that can be configured to save up to 4 MB. You can monitor system messages by clicking the **Events** tab on Fabric Manager or by choosing **Logs > Events > Current** on Device Manager. You can also monitor system messages remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a syslog server.



### Note

When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a syslog server for a few seconds.

## Syslog Server Logging Facilities and Severity Levels

All syslog messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single syslog daemon (syslogd) sends the information based on the configured facility option. If no facility is specified, local7 is the default outgoing facility.

The outgoing logging facilities are listed in [Table 31-1](#).

**Table 31-1**      **Outgoing Logging Facilities**

Facility Keyword	Description	Standard or Cisco MDS Specific
<b>auth</b>	Authorization system	Standard
<b>authpriv</b>	Authorization (private) system	Standard
<b>cron</b>	Cron or at facility	Standard
<b>daemon</b>	System daemons	Standard
<b>ftp</b>	File Transfer Protocol	Standard

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 31-1**      ***Outgoing Logging Facilities (continued)***

Facility Keyword	Description	Standard or Cisco MDS Specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal syslog messages	Standard
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard

Table 31-2 describes some samples of internal facilities supported by the system message logs.

**Table 31-2**      ***Internal Logging Facilities***

Facility Keyword	Description	Standard or Cisco MDS Specific
acl	ACL manager	Cisco MDS 9000 Family specific
all	All facilities	Cisco MDS 9000 Family specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
bootvar	Bootvar	Cisco MDS 9000 Family specific
callhome	Call Home	Cisco MDS 9000 Family specific
cron	Cron or at facility	Standard
daemon	System daemons	Standard
fcc	FCC	Cisco MDS 9000 Family specific
fcdomain	fcdomain	Cisco MDS 9000 Family specific
fcns	Name server	Cisco MDS 9000 Family specific
fcs	FCS	Cisco MDS 9000 Family specific
flogi	FLOGI	Cisco MDS 9000 Family specific
fspf	FSPF	Cisco MDS 9000 Family specific
ftp	File Transfer Protocol	Standard
ipconf	IP configuration	Cisco MDS 9000 Family specific
ipfc	PFC	Cisco MDS 9000 Family specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard
lpr	Line printer system	Standard
mail	Mail system	Standard
mcast	Multicast	Cisco MDS 9000 Family specific
module	Switching module	Cisco MDS 9000 Family specific

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 31-2 Internal Logging Facilities (continued)**

Facility Keyword	Description	Standard or Cisco MDS Specific
<b>news</b>	USENET news	Standard
<b>ntp</b>	NTP	Cisco MDS 9000 Family specific
<b>platform</b>	Platform manager	Cisco MDS 9000 Family specific
<b>port</b>	Port	Cisco MDS 9000 Family specific
<b>port-channel</b>	PortChannel	Cisco MDS 9000 Family specific
<b>qos</b>	QoS	Cisco MDS 9000 Family specific
<b>rdl</b>	RDL	Cisco MDS 9000 Family specific
<b>rib</b>	RIB	Cisco MDS 9000 Family specific
<b>rscn</b>	RSCN	Cisco MDS 9000 Family specific
<b>securityd</b>	Security	Cisco MDS 9000 Family specific
<b>syslog</b>	Internal system messages	Standard
<b>sysmgr</b>	System manager	Cisco MDS 9000 Family specific
<b>tlport</b>	TL port	Cisco MDS 9000 Family specific
<b>user</b>	User process	Standard
<b>uucp</b>	UNIX-to-UNIX Copy Program	Standard
<b>vhbad</b>	Virtual host base adapter daemon	Cisco MDS 9000 Family specific
<b>vni</b>	Virtual network interface	Cisco MDS 9000 Family specific
<b>vrrp_cfg</b>	VRRP configuration	Cisco MDS 9000 Family specific
<b>vrrp_eng</b>	VRRP engine	Cisco MDS 9000 Family specific
<b>vsan</b>	VSAN system messages	Cisco MDS 9000 Family specific
<b>vshd</b>	vshd	Cisco MDS 9000 Family specific
<b>wwn</b>	WWN manager	Cisco MDS 9000 Family specific
<b>xbar</b>	Xbar system messages	Cisco MDS 9000 Family specific
<b>zone</b>	Zone server	Cisco MDS 9000 Family specific

Table 31-3 describes the severity levels supported by the system message logs.

**Table 31-3 Error Message Severity Levels**

Level Keyword	Level	Description	System Message Definition
<b>emergency</b>	0	System unusable	LOG_EMERG
<b>alert</b>	1	Immediate action needed	LOG_ALERT
<b>critical</b>	2	Critical conditions	LOG_CRIT
<b>error</b>	3	Error conditions	LOG_ERR
<b>warning</b>	4	Warning conditions	LOG_WARNING
<b>notify</b>	5	Normal but significant condition	LOG_NOTICE



**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Table 31-3 Error Message Severity Levels (continued)**

Level Keyword	Level	Description	System Message Definition
info	6	Informational messages only	LOG_INFO
debug	7	Debugging messages	LOG_DEBUG



**Note**

Refer to the *Cisco MDS 9000 Family System Messages Reference* for details on the error log message format.

## Configuring Message Logging

System logging messages are sent to Fabric Manager or Device Manager based on the default (or configured) logging facility and severity values. You can enable logging to the console, terminal sessions, log file, or line cards using Fabric Manager or Device Manager.

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is critical.



**Note**

Most tabs in the Information pane for features using CFS are dimmed until you click the CFS tab. The CFS tab shows which switches have CFS enabled and shows the master switch for this feature. Once the CFS tab is click, the other tabs in the Information pane that use CFS are activated.



**Tip**

The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generate an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud

Logging messages may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages. The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes. The configured log file is saved in the /var/log/external directory. The location of the log file cannot be changed.

By default, logging is enabled at the debug level for all line cards. You can enable or disable logging for each line card at a specified level.

To enable or disable message logging for these features, follow these steps:

- Step 1** In Fabric Manager, choose **Switches > Events > Syslog** and click the **Switch Logging** tab in the Information pane.  
In Device Manager, choose **Logs > Syslog > Setup** and click the **Switch Logging** tab in the Syslog dialog box.
- Step 2** Check the check boxes for where you want message logging to occur.
- Step 3** Choose the message severity threshold from the **MsgSeverity** drop-down box for each switch in Fabric Manager, or click the appropriate message severity level radio button in Device Manager.
- Step 4** Click **Apply Changes** on Fabric Manager, or **Apply** on Device Manager to save and apply your changes.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring a Syslog Server

You can configure a maximum of three syslog servers. One of these syslog servers should be Fabric Manager if you want to view system messages from the Event tab in Fabric Manager.

To configure syslog servers, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | In Fabric Manager, choose <b>Switches &gt; Events &gt; Syslog</b> and click the <b>Servers</b> tab in the Information pane.<br><br>In Device Manager, choose <b>Logs &gt; Syslog &gt; Setup</b> and click the <b>Servers</b> tab in the Syslog dialog box. |
| <b>Step 2</b> | Click <b>Create Row</b> on Fabric Manager, or <b>Create</b> on Device Manager to add a new syslog server.  |
| <b>Step 3</b> | Enter the name or IP address in dotted decimal notation (for example, 192.168.2.12) of the syslog server in the Name or IP Address field.  |
| <b>Step 4</b> | Set the message severity threshold by clicking the <b>MsgSeverity</b> radio button and set the facility by clicking the Facility radio button.   |
| <b>Step 5</b> | Click <b>Apply Changes</b> on Fabric Manager, or click <b>Create</b> on Device Manager to save and apply your changes.   |
| <b>Step 6</b> | If CFS is enabled on Fabric Manager for the syslog feature, click <b>CFS</b> and commit these changes to propagate the configuration through the fabric.   |
- 

Device Manager allows you to view event logs on your local PC as well as those on the switch. For a permanent record of all events that occur on the switch, you should store these messages off the switch. To do this the MDS switch must be configured to send syslog messages to your local PC and a syslog server must be running on that PC to receive those messages. These messages can be categorized into four classes:

- Hardware—Line card or power supply problems
- Link Incidents—FICON port condition changes
- Accounting—User change events
- Events—All other events



### Note

You should avoid using PCs that have IP addresses randomly assigned to them by DHCP. The switch continues to use the old IP address unless you manually change it; however the Device Manager prompts you if it does detect this situation. UNIX workstations have a built-in syslog server. You must have root access (or run the Cisco syslog server as setuid to root) to stop the built-in syslog daemon and start the Cisco syslog server.

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Verifying Syslog Servers from Fabric Manager Web Services

To verify the syslog servers remotely from Fabric Manager Web Services, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Point your browser at the Fabric Manager Web Services server. See the <a href="#">“Launching and Using Fabric Manager Web Services” section on page 5-7</a> . |
| <b>Step 2</b> | Choose <b>Admin &gt; Events</b> to view the syslog server information for each switch. The columns in the table are sortable.                                 |
- 

## Viewing Logs from Fabric Manager Web Services

To view system messages remotely from Fabric Manager Web Services, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Point your browser at the Fabric Manager Web Services server. See the <a href="#">“Launching and Using Fabric Manager Web Services” section on page 5-7</a> .  |
| <b>Step 2</b> | Choose <b>Events &gt; Details</b> to view the system messages. The columns in the events table are sortable. In addition, you can use the Filter button to limit the scope of messages within the table. |
- 

## Viewing Logs from Device Manager

You can view system messages from Device Manager if Device Manager is running from the same workstation as the Fabric Manager Server. Choose **Logs > Events > current** to view the system messages on Device Manager. The columns in the events table are sortable. In addition, you can use the Find button to locate text within the table.

You can view switch-resident logs even if you have not set up your local syslog server or your local PC is not in the switch's syslog server list. Due to memory constraints, these logs will wrap when they reach a certain size. The switch syslog has two logs: an NVRAM log that holds a limited number of critical and greater messages and a nonpersistent log that contains notice or greater severity messages. Hardware messages are part of these logs.

## Health and Event Monitoring

Fabric Manager works with the Cisco MDS 9000 Family switches to show the health and status of the fabric and switches. Information about the fabric and its components is gathered from multiple sources, including Online System Health Management, Call Home, system messages, and SNMP notifications. This information is then made available from multiple menus on Fabric Manager or Device Manager.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Fabric Manager Events Tab

The Fabric Manager Events tab, available from the topology window, displays the events Fabric Manager received from sources within the fabric. These sources include SNMP events, RMON events, system messages, and system health messages. The Events tab shows a table of events, including the event name, the source and time of the event, a severity level, and a description of the event. The table is sortable by any of these column headings.

## Event Information in Fabric Manager Web Services Reports

The Fabric Manager web services client displays collections of information gathered by the Performance Manager. This information includes events sent to the Fabric Manager Server from the fabric. To open these reports, choose **Performance Manager > Reports**. This opens the web client in a web browser and displays a summary of all fabrics monitored by the Fabric Manager Server. Choose a fabric and then click the **Events** tab to see a summary or detailed report of the events that have occurred in the selected fabric. The summary view shows how many switches, ISLs, hosts, or storage elements are down on the fabric and how many warnings have been logged for that SAN entity. The detailed view shows a list of all events that have been logged from the fabric and can be filtered by severity, time period, or type.

## Events in Device Manager

Device Manager displays the events when you choose **Logs > Events**. Device Manager can display the current list of events or an older list of events that has been stored on the Fabric Manager host. The event table shows details on each event, including time, source, severity, and a brief description of the event.



## Performance Monitoring

---

Cisco Fabric Manager and Device Manager provide multiple tools for monitoring the performance of the overall fabric, SAN elements, and SAN links. These tools provide real-time statistics as well as historical performance monitoring.

This section contains the following topics:

- [Real-Time Performance Monitoring, page 32-1](#)
- [Historical Performance Monitoring, page 32-2](#)

### Real-Time Performance Monitoring

Real-time performance statistics are a useful tool in dynamic troubleshooting and fault isolation within the fabric. Real-time statistics gather data on parts of the fabric in user-configurable intervals and display these results in Fabric Manager and Device Manager.

### Device Manager Real-Time Performance Monitoring

Device Manager provides an easy tool for monitoring ports on the Cisco MDS 9000 Family switches. This tool gathers statistics at a configurable interval and displays the results in tables or charts. These statistics show the performance of the selected port in real-time and can be used for performance monitoring and troubleshooting. For a selected port, you can monitor any of a number of statistics including traffic in and out, errors, class 2 traffic, and FICON data. You can set the polling interval from ten seconds to one hour, and display the results based on a number of selectable options including absolute value, value per second, and minimum or maximum value per second.

Device manager provides two performance views, the Summary View tab, and the configurable monitor option per port.

To configure the summary view in Device Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Click the <b>Summary</b> tab on the main display. You see all of the active ports on the switch, as well as the configuration options available from the Summary view.  |
| <b>Step 2</b> | Choose the <b>Poll Interval</b> and how the data should be interpreted by clicking the <b>Show Rx/Tx</b> drop-down menu. The table updates each polling interval to show an overview of the receive and transmit data for each active port on the switch. |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 3** Choose **Show Rx/Tx > % Util/sec**, and then choose the warning and critical threshold levels for event reporting. You can also display the percent utilization for a port by selecting the port and clicking the **Monitor Selected Interface Traffic Util %** icon.
- Step 4** Click the **Monitor Selected Interface Traffic Details** icon if you want more detailed statistics on the port.

---

The configurable monitor option per port gives statistics for in and out traffic on that port, errors, class 2 traffic and other data that can be graphed over a period of time to give a real-time view into the performance of the port.

To configure per port monitoring in Device Manager, follow these steps:

- 
- Step 1** Right-click the port you are interested in and choose **Monitor...** from the options pop-up menu. You see the port real-time monitor dialog box.
- Step 2** Choose the **Poll Interval** and how the data should be interpreted using the drop-down menu. The table updates each polling interval to show statistics for the selected port.
- Step 3** Click a statistics value from the table and then click one of the graphing icons to display a running graph of that statistic over time. You see a graph window that contains options to change the graph type.



**Tip** You can open multiple graphs for statistics on any of the active ports on the switch.

---

## Fabric Manager Real-Time ISL Statistics

You can configure Fabric Manager to gather ISL statistics in real time. These ISL statistics include receive and transmit utilization, bytes per second, as well as errors and discards per ISL.

To configure ISL statistics in Fabric Manager, follow these steps:

- 
- Step 1** Select **Performance > ISL in Real-Time**. ISL statistics display in the Information pane.
- Step 2** Choose the **Poll Interval** and bandwidth utilization thresholds. The table updates each polling interval to show statistics for all configured ISLs in the fabric.
- Step 3** Select a row in the table to highlight that ISL in blue on the Topology map.
- 

## Historical Performance Monitoring

Performance Manager gathers network device statistics historically and provides this information graphically using a web browser. It presents recent statistics in detail and older statistics in summary. Performance Manager also integrates with external tools such as Cisco Traffic Analyzer.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

The Performance Manager has three operational stages:

- Definition—Uses two configuration wizards to create a collection configuration file.
- Collection—Reads the configuration file and collects the desired information.
- Presentation—Generates web pages to present the collected data.

See the “[Performance Manager Architecture](#)” section on page 6-1 for an overview of Performance Manager.

## Creating a Flow with Performance Manager

Performance Manager has a Flow Configuration Wizard that steps you through the process of creating host-to-storage, storage-to-host, or bidirectional flows. [Table 32-1](#) explains the Flow Type radio button that defines the type of traffic monitored.

**Table 32-1** Performance Manager Flow Types

Flow type	Description
Host->Storage	Unidirectional flow, monitoring data from the host to the storage element
Storage->Host	Unidirectional flow, monitoring data from the storage element to the host
Both	Bidirectional flow, monitoring data to and from the host and storage elements.

Once defined, these flows can be added to a collection configuration file to monitor the traffic between a host/storage element pair.

To create a flow in Fabric Manager, follow these steps:

- 
- Step 1** Choose **Performance > Create Flows** to launch the wizard.
  - Step 2** Choose the **VSAN** from which you want to create flows. Flows are defined per VSAN.
  - Step 3** Click the **Type** radio button for the flow type you want to define.
  - Step 4** Check the **Clear old flows on modified switches** check box if you want to remove old flow data.
  - Step 5** Click **Next** to review the available flows for the chosen VSAN. Remove any flows you are not interested in.
  - Step 6** Click **Finish** to create the flow.
- 

The flows created become part of the collection options in the Performance Manager Configuration Wizard.

## Creating a Collection with Performance Manager

The Performance Manager Configuration Wizard steps you through the process of creating collections using configuration files. Collections are defined for one or all VSANs in the fabric. Collections can include statistics from the SAN element types described in [Table 32-2](#).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 32-2 Performance Manager Collection Types**

Collection Type	Description
ISLs	Collects link statistics for ISLs.
Host	Collects link statistics for SAN hosts.
Storage	Collects link statistics for a storage elements.
Flows	Collects flow statistics defined by the Flow Configuration Wizard.

## Using Performance Thresholds

The Performance Manager Configuration Wizard allows you to set up two thresholds that will trigger events when the monitored traffic exceeds the percent utilization configured. These event triggers can be set as either Critical or Warning events that are reported on the Fabric Manager web client Events browser page.

You must choose either absolute value thresholds or baseline thresholds that apply to all transmit or receive traffic defined in the collection. Click the **Use absolute values** radio button on the last screen of the Performance Manager Configuration Wizard to configure thresholds that apply directly to the statistics gathered. These statistics, as a percent of the total link capacity, are compared to the percent utilization configured for the threshold type. If the statistics exceed either configured threshold, an event is shown on the Fabric Manager web client Events tab.

As an example, the collection has absolute value thresholds set for 60% utilization (for warning) and 80% utilization (for critical). If Performance Manager detects that the traffic on a 1-Gigabit link in its collection exceeds 600 Mbps, a warning event is triggered. If the traffic exceeds 800 Mbps, a critical event is triggered.

Baseline thresholds are defined for a configured time of day or week (1 day, 1 week, or 2 weeks). The baseline is created by calculating the average of the statistical results for the configured time each day, week, or every 2 weeks. [Table 32-3](#) shows an example of the statistics used to create the baseline value for a collection defined at 4 pm on a Wednesday.

**Table 32-3 Baseline Time Periods for a Collection Started on Wednesday at 4pm**

Baseline Time Window	Statistics Used in Average Calculation
1 day	Every prior day at 4 pm
1 week	Every prior Wednesday at 4 pm
2 weeks	Every other prior Wednesday at 4 pm

Baseline thresholds create a threshold that adapts to the typical traffic pattern for each link for the same time window each day, week, or every 2 weeks. Baseline thresholds are set as a percent of the average (110% to 500%), where 100% equals the calculated average.

As an example, a collection is created at 4 pm on Wednesday, with baseline thresholds set for 1 week, at 150% of the average (warning) and 200% of the average (critical). Performance Manager recalculates the average for each link at 4 pm every Wednesday by taking the statistics gathered at that time each Wednesday since the collection started. Using this as the new average, Performance Manager compares each received traffic statistic against this value and sends a warning or critical event if the traffic on a link exceeds this average by 150% or 200% respectively.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Table 32-4 shows two examples of 1-Gigabit links with different averages in our example collection and at what traffic measurements the Warning and Critical events are sent.

**Table 32-4 Example of Events Generated for 1-Gigabit Links**

Average	Warning Event Sent at 150%	Critical Event Sent at 200%
400 Mbps	600 Mbps	800 Mbps
200 Mbps	300 Mbps	400 Mbps

Set these thresholds on the last screen of the Collections Configuration Wizard by checking the **Send events if traffic exceeds threshold** check box.

## Using the Performance Manager Configuration Wizard

To create a collection using the Performance Manager Configuration Wizard in Fabric Manager, follow these steps:

- Step 1** Choose **Performance > Create Collection** to launch the Performance Manager Configuration Wizard.
- Step 2** Choose the VSANs from which you want to collect data or choose **All** to collect statistics across all VSANs in the fabric.
- Step 3** Check the **Type** check boxes for each type of links flow or SAN element that you want included in your collection.
- Step 4** If you want to ignore flows with zero counter values, check that check box.
- Step 5** If you are using Cisco Traffic Analyzer, enter the URL where it is located on your network.
- Step 6** Click **Next** to review the collection specification data. Remove any links, flows, or SAN elements you are not interested in.
- Step 7** Click **Next** to configure other collection options.
- Step 8** Check the appropriate check boxes if you want to include errors and discards in your collection, and if you want to interpolate data for missing statistics.
- Step 9** Check the **Send event if traffic exceeds threshold** check box if you want to configure threshold events as explained in the [“Using Performance Thresholds”](#) section on page 32-4.
- Step 10** Click the **Use absolute values** radio button if you want absolute value thresholds or click the **Baseline values over** radio button if you want baseline thresholds.
- Step 11** Choose the time window for baseline calculations if baseline thresholds are configured.
- Step 12** Choose the Critical and Warning threshold values as a percent of link capacity (for absolute value thresholds ) or average (for baseline thresholds).
- Step 13** Click **Finish** to create the collection configuration file. You see a dialog box asking if you want to restart Performance Manager.
- Step 14** Click **Yes** to restart Performance Manager to use this new configuration file, or click **No** to exit the Performance Manager Configuration Wizard without restarting Performance Manager. If you choose No, Performance Manager will not use the new configuration file until you restart it by choosing **Performance Manager > Collector > Restart**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Note**

If you reconfigure your fabric, you may need to update your Performance Manager collections and flows. Recreate your flows and collections using Performance Manager Configuration Wizard.

## Starting and Stopping Data Collection

After configuring the collection or iSCSI flows, you can start or restart the collection by choosing **Performance > Collector > Start** or **Performance > Collector > Restart**. You can verify that the collection has started by checking the PMCollector.log file in the main MDS9000 directory on your Fabric Manager Server, or by viewing the status of the collection on the Fabric Manager web client Admin tab.

You can manually stop a data collection process in Windows using the services panel. Right-click the **Cisco Performance Manager** service and choose **Stop**.

On a UNIX machine, enter the following command:

```
$HOME/.ciscomds9000/bin/pm.sh stop
```

You can also start, restart, or stop the collection on the Fabric Manager web client Admin tab.

## Viewing Performance Manager Reports

You can view Performance Manager statistical data using preconfigured reports that are built on demand and displayed in a web browser. These reports provide summary information as well as detailed statistics that can be viewed for daily, weekly, monthly, or yearly results.

Choose **Performance > Reports** to access Performance Manager reports from Fabric Manager. This opens a web browser window showing the default Fabric Manager web client event summary report. Click the **Performance** tab to view the Performance Manager reports. Performance Manager begins reporting data ten minutes after the collection is started.

## Performance Summary

The Performance Summary page presents a dashboard display of the throughput and link utilization for hosts, ISLs, storage, and flows for the last 24-hour period. The summary provides a quick overview of the fabric's bandwidth consumption and highlights any hotspots.

The report includes network throughput pie charts and link utilization pie charts. Use the navigation tree on the left to show summary reports for monitored fabrics or VSANs. The summary displays charts for all hosts, storage elements, ISLs, and flows. Each pie chart shows the percent of entities (links, hosts, storage, ISLs, or flows) that measure throughput or link utilization on each of six predefined ranges. Move the mouse over a pie chart section to see how many entities exhibit that range of statistics. Double-click any pie chart to bring up a table of statistics for those hosts, storage elements, ISLs, or flows.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Performance Tables and Details Graphs

Click **Host**, **Storage**, **ISL**, or **Flow** to view traffic over the past day for all hosts, storage, ISLs, or flows respectively. A table lists all of the selected entities, showing transmit and receive traffic and errors and discards, if appropriate. The table can be sorted by any column heading. The table can also be filtered by day, week, month, or year. Tables for each category of statistics display average and peak throughput values and provide hot-links to more detailed information.

Clicking a link in any of the tables opens a details page that shows graphs for traffic by day, week, month, and year. If flows exist for that port, you can see which storage ports sent data. The details page also displays graphs for errors and discards if they are part of the statistics gathered and are not zero.

If you double-click a graph on a Detail report, it will launch the Cisco Traffic Analyzer for Fibre Channel, if configured. The aliases associated with hosts, storage devices, and VSANs in the fabric are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

## Viewing Performance Manager Events

Performance Manager events are viewable through the Fabric Manager Web Client.

To view Performance Manager events, follow these steps:

- 
- Step 1** Choose **Performance Manager > Reports**. You see a summary of all fabrics monitored by the Fabric Manager Server in a web browser.
  - Step 2** Choose a fabric and then click the **Events** tab to see a summary or detailed report of the events that have occurred in the selected fabric.
- 

## Generating Top10 Reports in Performance Manager

Cisco MDS SAN-OS Release 2.1(1a) introduces the ability to generate historical Top10 reports that can be saved for later review. These reports list the entities from the data collection, with the most active entities appearing first. This is a static, one-time only report that generates averages and graphs of the data collection as a snapshot at the time the report is generated.



### Tip

Name the reports with a timestamp so that you can easily find the report for a given day or week.

These Top10 reports differ from the other monitoring tables and graphs in Performance Manager in that the other data is continuously monitored and is sortable on any table column. The Top10 reports are a snapshot view at the time the report was generated.



### Note

Top10 reports require analyzing the existing data over an extended period of time and can take hours or more to generate on large fabrics.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To generate a Top10 report using Fabric Manager Web Services, follow these steps:

- 
- Step 1** Choose **Performance** from the main page and click the **Reports** tab. The list of existing reports displays.
  - Step 2** Enter a name for a new report and click **Generate**. The report may take hours to generate. When finished, the report appears by name in the left-hand navigation bar.
  - Step 3** Click the name of the generated report to see the Top10 tables for your fabric.
  - Step 4** Click the name of any entity in the Top10 tables to see a series of graphs for the transmit and receive data rates as well as errors and discards.
- 

## Generating Top10 Reports Using Scripts

You can generate Top10 reports manually by issuing the following commands:

- On UNIX, run the script:

```
"/<user_directory>/cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>"
```

- On Windows, run the script:

```
"c:\Program Files\Cisco Systems\MDS 9000\bin\pm.bat display pm\pm.xml
<output_directory>"
```

On UNIX, you can automate the generation of the Top10 reports on your Fabric Manager Server host by adding the following cron entry to generate the reports once an hour:

```
0 * * * * /<user_directory>/cisco_mds9000/bin/pm.sh display pm/pm.xml <output_directory>
```

If your crontab does not run automatically or Java complains about an exception similar to [Example 32-1](#), you need to add “-Djava.awt.headless=true” to the JVMARGS command in /<user\_directory>/cisco\_mds9000/bin/pm.sh.

### **Example 32-1 Example Java Exception**

```
in thread "main" java.lang.InternalError Can't connect to X11 window server using '0.0' as
the value of the DISPLAY variable.
```

## Exporting Data Collections to XML Files

The RRD files used by Performance Manager can be exported to a freeware tool called rrdtool. The rrd files are located in pm/db on the Fabric Manager Server. To export the collection to an XML file, enter the following command at the operating system command-line prompt:

```
/bin/pm.bat xport xxx yyy
```

In this command, xxx is the RRD file and yyy is the XML file that is generated. This XML file is in a format that rrdtool is capable of reading with the command:

```
rrdtool restore filename.xml filename.rrd
```

You can import an XML file with the command:

```
bin/pm.bat pm restore <xmlFile> <rrdFile>
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

This reads the XML export format that rrdtool is capable of writing with the command:

```
rrdtool xport filename.xml filename.rrd.
```

The **pm xport** and **pm restore** commands can be found on your Fabric Manager Server at bin\PM.bat for Windows platforms or bin/PM.sh on UNIX platforms. For more information on the rrdtool, refer to the following website: <http://www.rrdtool.org>.

## Exporting Data Collections in Readable Format

Cisco MDS SAN-OS Release 2.1(1a) introduces the ability to export data collections in comma-separated format (CSV). This format can be imported to various tools, including Microsoft Excel. You can export these readable data collections either from the Fabric Manager Web Services menus or in batch mode from the command line on Windows or UNIX. Using Fabric Manager Web Services, you can export one file. Using batch mode, you can export all collections in the pm.xml file.

To export data collections using Fabric Manager Web Services, follow these steps:

- 
- Step 1** Choose **Performance** from the main page and select the category you want to export from (Hosts, Storage, ISLs, or Flows). You see the overview table.
  - Step 2** Double-click the **Name** of the entity you want to export. You see the detailed graph for that entity in a pop-up window.
  - Step 3** Click the **Export** icon in the lower left of the graph. You see the Open/Save dialog box.
  - Step 4** Click **Save** to save the csv file or click **Cancel** to discard this operation.
- 

To export data collections using command line batch mode, follow these steps:

- 
- Step 1** Go to the installation directory on your workstation and then go to the bin directory.
  - Step 2** On Windows, enter **.pm.bat export C:\Program Files\Cisco Systems\MDS 9000\pm\pm.xml <export directory>**. This creates the csv file, export.csv, in the <export directory> on your workstation.
  - Step 3** On UNIX, enter **./pm.sh export /usr/local/cisco\_mds9000/pm/pm.xml <export directory>**. This creates the csv file, export.csv, in the <export directory> on your workstation.
- 

When you open this exported file in Microsoft Excel, the following information displays:

- Title of the entity you exported and the address of the switch the information came from.
- The maximum speed seen on the link to or from this entity.
- The VSAN ID and maximum speed.
- The timestamp, followed by the receive and transmit data rates in bytes per second.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring Performance Manager for Use with Cisco Traffic Analyzer

Performance Manager works in conjunction with the Cisco Traffic Analyzer to allow you to monitor and manage the traffic on your fabric. Using Cisco Traffic Analyzer with Performance Manager requires the following components:

- A configured Fibre Channel Switched Port Analyzer (SPAN) destination (SD) port to forward Fibre Channel traffic.
- A Port Analyzer Adapter 2 (PAA-2) to convert the Fibre Channel traffic to Ethernet traffic.
- Cisco Traffic Analyzer software to analyze the traffic from the PAA-2.

To configure Performance Manager to work with the Cisco Traffic Analyzer, follow these steps:

- 
- Step 1** Set up the Cisco Traffic Analyzer according to the instructions in the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.
- Step 2** Get the following three pieces of information:
- The IP address of the management workstation on which you are running Performance Manager and Cisco Traffic Analyzer.
  - The path to the directory where Cisco Traffic Analyzer is installed.
  - The port that is used by Cisco Traffic Analyzer (the default is 3000).
- Step 3** Start the Cisco Traffic Analyzer.
- a. Choose **Performance > Traffic Analyzer > Open**.
  - b. Enter the URL for the Cisco Traffic Analyzer, in the format  
`http://<ip address>:<port number>`  
where:  
<ip address> is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and  
:<port number> is the port that is used by Cisco Traffic Analyzer (the default is :3000).
  - c. Click **OK**.
  - d. Choose **Performance > Traffic Analyzer > Start**.
  - e. Enter the location of the Cisco Traffic Analyzer, in the format  
`D:\<directory>\ntop.bat`  
where:  
D: is the drive letter for the disk drive where the Cisco Traffic Analyzer is installed, and  
<directory> is the directory containing the ntop.bat file.
  - f. Click **OK**.
- Step 4** Create the flows you want Performance Manager to monitor, using the Flow Configuration Wizard.
- Step 5** Define the data collection you want Performance Manager to gather, using the Performance Manager Configuration Wizard.
- a. Choose the VSAN you want to collect information for or choose **All VSANs**.
  - b. Check the types of items you want to collect information for (hosts, ISLs, storage devices, and flows).

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

- c. Enter the URL for the Cisco Traffic Analyzer in the format

`http://<ip address>/<directory>`

where:

`<ip address>` is the address of the management workstation on which you have installed the Cisco Traffic Analyzer, and `<directory>` is the path to the directory where the Cisco Traffic Analyzer is installed.

- d. Click **Next**.
- e. Review the data collection on this and the next section to make sure this is the data you want to collect.
- f. Click **Finish** to begin collecting data.



**Note**

Data is not collected for JBOD or for virtual ports. If you change the data collection configuration parameters during a data collection, you must stop and restart the collection process for your changes to take effect.

- Step 6** Choose **Performance > Reports** to generate a report.



**Note**

It takes at least five minutes to start collecting data for a report. Do not attempt to generate a report in Performance Manager during the first five minutes of collection.

- Step 7** Click the **Cisco Traffic Analyzer** at the top of the Host or Storage detail pages to view the Cisco Traffic Analyzer information, or choose **Performance > Traffic Analyzer > Open**. The Cisco Traffic Analyzer page will not open unless ntop has been started already.



**Note**

For information on capturing a SPAN session and starting a Cisco Traffic Analyzer session to view it, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.



**Note**

For information on viewing and interpreting your Performance Manager data, see the [“Historical Performance Monitoring” section on page 32-2](#).

For information on viewing and interpreting your Cisco Traffic Analyzer data, refer to the *Cisco MDS 9000 Family Port Analyzer Adapter 2 Installation and Configuration Note*.

For performance drill-down, Fabric Manager Server can launch the Cisco Traffic Analyzer in-context from the Performance Manager graphs. The aliases associated with hosts, storage devices, and VSANs are passed to the Cisco Traffic Analyzer to provide consistent, easy identification.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Third-Party Integration

---

Fabric Manager provides tools to facilitate integration with third-party applications and devices. This chapter contains the following sections:

- [Call Home Configuration, page 33-1](#)
- [Configuring SNMP Events, page 33-11](#)
- [Configuring RMON Using Threshold Manager, page 33-13](#)

### Call Home Configuration

The actual configuration of Call Home depends on how you intend to use the feature. Some points to consider include:

- An e-mail server and at least one destination profile (predefined or user-defined) must be configured. The destination profile(s) used depends on whether the receiving entity is a pager, e-mail, or automated service such as Cisco AutoNotify.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This is required to determine the origin of messages received.
- The Cisco MDS 9000 switch must have IP connectivity to an e-mail server.
- You must obtain an active service contract to use Cisco AutoNotify.

### Cisco AutoNotify

If you have service contracts directly with Cisco Systems, automatic case generation with the Technical Assistance Center is possible by registering with the AutoNotify service. AutoNotify provides fast time to resolution of system problems by providing a direct notification path to Cisco customer support.

The AutoNotify feature requires several Call Home parameters to be configured, including certain contact information, e-mail server, and an XML destination profile as specified in the Service Activation document found on the Cisco.com website at:

[http://www.cisco.com/univercd/cc/td/doc/product/voice/c\\_callmg/3\\_3/service/serv332/ccmsrvs/sssrvac t.htm](http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/3_3/service/serv332/ccmsrvs/sssrvac t.htm)

To configure a Cisco MDS 9000 Family switch to use the AutoNotify service, an XML destination profile must be configured to send messages to Cisco. Specific setup, activation, and e-mail address information is found on the Cisco.com website at:

[http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti\\_ds.htm](http://www.cisco.com/warp/customer/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti_ds.htm)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To register, the following items are required:

- The SMARTnet contract number covering your Cisco MDS 9000 Family switch.
- Your name, company address, e-mail address, and Cisco.com ID.
- The exact product number of your Cisco MDS 9000 Family switch. For example, some valid product numbers include: DS-C6509 and DS-C9216-K9.
- The serial number of your Cisco MDS 9000 Family switch. This can be obtained by looking at the serial number label on the back of the switch (next to the power supply).

The ContractID, CustomerID, SiteID, and SwitchPriority parameters are not required by the AutoNotify feature. They are only intended to be used as additional information by Cisco customers and service partners.

## Configuring Call Home

To configure Call Home, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Expand the <b>Switches</b> folder in the Physical Attributes pane on Fabric Manager and choose <b>Events &gt; Call Home</b> . You see the Call Home dialog box in the Information pane.<br>Or, choose <b>Admin &gt; Events &gt; Call Home</b> on Device Manager. |
| <b>Step 2</b> | Click the <b>General</b> tab to assign contact information and enable the Call Home feature. The Call Home feature is not enabled by default, and you must enter an e-mail address that identifies the source of Call Home notifications.                        |
| <b>Step 3</b> | Click the <b b="" destinations<=""> tab to configure the destination e-mail addresses for Call Home notifications. You can identify one or more e-mail addresses that will receive Call Home notifications.</b>  |
| <b>Step 4</b> | Click the <b>E-mail Setup</b> tab to identify the SMTP server. You need to identify a message server to which your switch has access. This message server will forward the Call Home notifications to the destinations.  |
- 

## Configuring Call Home Destination Profiles and Alert Groups

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. At least one destination profile is required. You can configure multiple destination profiles of one or more types.



### Note

If you use the Cisco AutoNotify service, the XML destination profile is required (see [http://www.cisco.com/warp/public/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti\\_ds.htm](http://www.cisco.com/warp/public/cc/serv/mkt/sup/tsssv/opmsup/smtan/anoti_ds.htm)).

A destination profile consists of the following:

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- **Message severity**—Severity of messages that will trigger a Call Home message. Messages below this severity do not trigger Call Home messages.
- **Alert group**—A predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Family

Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

To configure Call Home profiles from the Fabric Manager, follow these steps:

- 
- Step 1** Choose **Switches > Events > Call Home** from the Physical Attributes tree and click the **Profiles** tab in the Information pane. You see Call Home profiles for multiple switches.
- Step 2** Click the **Create Row** icon to add a new profile.
- Step 3** Set the profile name, message format, size, and severity level.
- Step 4** Check the check boxes for each alert group you want sent in this profile.
- Step 5** Click **Create** to create this profile on the selected switches.
- 

## Call Home Message Severity Levels

You can filter Call Home messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower than the urgency threshold is not sent. The urgency level ranges from debug (lowest level of urgency) to catastrophic (highest level of urgency), and the default is debug (all messages are sent).



### Note

Call Home severity levels are not the same as system message logging severity levels.

To set the message level for each profile for Call Home, follow these steps:

- 
- Step 1** Expand the **Switches** folder in the Physical Attributes pane on Fabric Manager and choose **Events > Call Home**. You see the Call Home dialog box in the Information pane.  
Or, choose **Admin > Events > Call Home** on Device Manager.
- Step 2** Click the **Profiles** tab and set the message level for each switch using the drop-down menu in the **MsgLevel** column.
- Step 3** Click **Apply Changes** to save your changes or click **Undo Changes** to cancel your changes.
- 

## Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned commands to execute when the event occurs. The command output is included in the transmitted message. [Table 33-1](#) lists the trigger events.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-1 Event Triggers**

Event	Alert Group	Event Name	Description	Severity Level
Call Home	System and CISCO_TAC	SW_CRASH	A software process has crashed with a stateless restart, indicating an interruption of a service.	major
	System and CISCO_TAC	SW_SYSTEM_INCONSISTENT	Inconsistency detected in software or file system.	major
	Environmental and CISCO_TAC	TEMPERATURE_ALARM	Thermal sensor indicates temperature reached operating threshold.	critical
		POWER_SUPPLY_FAILURE	Power supply failed.	critical
		FAN_FAILURE	Cooling fan has failed.	major
	Switching module and CISCO_TAC	LINECARD_FAILURE	Switching module operation failed.	fatal
		POWER_UP_DIAGNOSTICS_FAILURE	Switching module failed power-up diagnostics.	fatal
	Line Card Hardware and CISCO_TAC	PORT_FAILURE	Hardware failure of interface port(s).	critical
	Line Card Hardware, Supervisor Hardware, and CISCO_TAC	BOOTFLASH_FAILURE	Failure of boot compact Flash card.	critical
	Supervisor module and CISCO_TAC	SUP_FAILURE	Supervisor module operation failed.	fatal
		POWER_UP_DIAGNOSTICS_FAILURE	Supervisor module failed power-up diagnostics.	fatal
Call Home	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of in-band communications path.	fatal
	Supervisor Hardware and CISCO_TAC	EOBC_FAILURE	Ethernet out-of-band channel communications failure.	critical
	Supervisor Hardware and CISCO_TAC	MGMT_PORT_FAILURE	Hardware failure of management Ethernet port.	major
	License	LICENSE_VIOLATION	Feature in use is not licensed (Cisco MDS SAN-OS Release 1.3.x), and is turned off after grace period expiration.	critical

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-1**      **Event Triggers (continued)**

Event	Alert Group	Event Name	Description	Severity Level
Inventory	Inventory and CISCO_TAC	COLD_BOOT	Switch is powered up and reset to a cold boot sequence.	notification
		HARDWARE_INSERTION	New piece of hardware inserted into the chassis.	notification
		HARDWARE_REMOVAL	Hardware removed from the chassis.	notification
Test	Test and CISCO_TAC	TEST	User generated test.	notification
Port syslog	Syslog-group-port	SYSLOG_ALERT	Syslog messages corresponding to the port facility.	notification
RMON	RMON	RMON_ALERT	RMON alert trigger messages.	notification

## Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- E-mail address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

[Table 33-2](#) describes the short text formatting option for all message types.

**Table 33-2**      **Short Text Messages**

Data Item	Description
Device identification	Configured device name
Date/time stamp	Time stamp of the triggering event
Error isolation message	Plain English description of triggering event
Alarm urgency level	Error level such as that applied to system message

## Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).

Table 33-3, Table 33-4, and Table 33-5 display the information contained in plain text and XML messages.

**Table 33-3**      **Reactive Event Message Format**

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .  <b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message.	/mml/header/name
Message type	Specifically “Call Home.”	/mml/header/type
Message group	Specifically “reactive.”	/mml/header/group
Severity level	Severity level of message.	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678”	/mml/ header/deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header/customerID
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header/siteId
Server ID	If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678”	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event. This is the host name of the device.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-3** *Reactive Event Message Format (continued)*

<b>Data Item (Plain text and XML)</b>	<b>Description (Plain text and XML)</b>	<b>XML Tag (XML only)</b>
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhone Number
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
Affected FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
Affected FRU serial number	Serial number of affected FRU.	/mml/body/fru/serialNo
Affected FRU part number	Part number of affected FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU generating the event message.	/mml/body/fru/slot
FRU hardware version	Hardware version of affected FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on affected FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment/ name
Attachment type	Specifically command output.	/mml/attachments/attachment/ type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/ mime
Command output text	Output of command automatically executed.	/mml/attachments/attachment/ atdata

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-4 Inventory Event Message Format**

Data Item (Plain text and XML)	Description (Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .  <b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically “Inventory Update.”	/mml/header/name
Message type	Specifically “Inventory Update.”	/mml/header/type
Message group	Specifically “proactive.”	/mml/header/group
Severity level	Severity level of inventory event is level 2.	/mml/header/level
Source ID	Product type for routing at Cisco. Specifically “MDS 9000.”	/mml/header/source
Device ID	Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /customerID
Contract ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field, used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Node that experienced the event.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-4** *Inventory Event Message Format (continued)*

<b>Data Item (Plain text and XML)</b>	<b>Description (Plain text and XML)</b>	<b>XML Tag (XML only)</b>
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the unit. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo
Chassis part number	Top assembly number of the chassis.	/mml/body/chassis/partNo
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion
Supervisor module software version	Top level software version.	/mml/body/chassis/swVersion
FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name
FRU s/n	Serial number of FRU.	/mml/body/fru/serialNo
FRU part number	Part number of FRU.	/mml/body/fru/partNo
FRU slot	Slot number of FRU.	/mml/body/fru/slot
FRU hardware version	Hardware version of FRU.	/mml/body/fru/hwVersion
FRU software version	Software version(s) running on FRU.	/mml/body/fru/swVersion
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name
Attachment type	Specifically command output.	/mml/attachments/attachment/type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime
Command output text	Output of command automatically executed after event categories (see ).	/mml/attachments/attachment/atdata

**Table 33-5** *User-Generated Test Message Format*

<b>Data Item (Plain text and XML)</b>	<b>Description (Plain text and XML)</b>	<b>XML Tag (XML only)</b>
Time stamp	Date and time stamp of event in ISO time notation: <i>YYYY-MM-DDTHH:MM:SS</i> .  <b>Note</b> The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time
Message name	Name of message. Specifically test message for test type message.	/mml/header/name
Message type	Specifically “Test Call Home.”	/mml/header/type

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-5**      **User-Generated Test Message Format (continued)**

<b>Data Item (Plain text and XML)</b>	<b>Description (Plain text and XML)</b>	<b>XML Tag (XML only)</b>
Message group	This field should be ignored by the receiving Call Home processing application, but may be populated with either “proactive” or “reactive.”	/mml/header/group
Severity level	Severity level of message, test Call Home message.	/mml/header/level
Source ID	Product type for routing.	/mml/header/source
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /customerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/ header /contractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/ header /siteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format: type@Sid@serial, where <ul style="list-style-type: none"> <li>Type is the product model number from backplane SEEPROM.</li> <li>@ is a separator character.</li> <li>Sid is “C” identifying serial ID as a chassis serial number.</li> <li>Serial: The serial number as identified by the Sid field.</li> </ul> Example: “DS-C9000@C@12345678	/mml/header/serverId
Message description	Short text describing the error.	/mml/body/msgDesc
Device name	Switch that experienced the event.	/mml/body/sysName
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact
Contact e-mail	E-mail address of person identified as contact for this unit.	/mml/body/sysContactEmail
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress
Model name	Model name of the switch. This is the specific model as part of a product family name.	/mml/body/chassis/name
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 33-5**      *User-Generated Test Message Format (continued)*

<b>Data Item (Plain text and XML)</b>	<b>Description (Plain text and XML)</b>	<b>XML Tag (XML only)</b>
Chassis part number	Top assembly number of the chassis. For example, 800-xxx-xxxx.	/mml/body/chassis/partNo
Command output text	Output of command automatically executed after event categories listed in .	/mml/attachments/attachment/atdata
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime
Attachment type	Specifically command output.	/mml/attachments/attachment/type
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name

## Configuring SNMP Events

SNMP events are asynchronous notifications of status, performance, or configuration changes on the monitored switch. These events can be either traps or informs. Traps are unacknowledged, while informs use TCP to tell the sending switch that the event was received at the configured destination. The switch will retry an inform after the configured timeout period if the SNMP event destination does not return an acknowledgement of the SNMP event. Fabric Manager or Device Manager can enable or disable individual SNMP events for each switch to provide customized notifications.

## Filtering SNMP Events

SNMP events cover a broad spectrum of features on the Cisco MDS 9000 Family switches. As an administrator, you may find some events are more important to your fabric management than others. Fabric Manager and Device Manager let you enable or disable individual SNMP events so that you only receive the SNMP events you are interested in at your SNMP event destinations.

To filter individual SNMP events, follow these steps:

- 
- Step 1**      On Fabric Manager, choose **Switches > Events > SNMP Traps** and then click the **FC** or **Other** tab in the Information pane.  
On Device Manager, choose **Admin > Events > Filters**.
  - Step 2**      Check the check boxes for the SNMP events you want sent to your SNMP event destinations.
  - Step 3**      Click **Apply Changes** on Fabric Manager, or **Create** on Device Manager to save and apply your changes.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring SNMP Event Destinations

Cisco MDS 9000 Family switches, like other SNMP-enabled devices, send events (traps and informs) to configured destinations, called *trap receivers* in SNMPv2.

To configure SNMP event destinations, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | On Fabric Manager, choose <b>Switches &gt; Events &gt; SNMP Traps</b> and then click the <b>Destinations</b> tab in the Information pane.<br><br>On Device Manager, choose <b>Admin &gt; Events &gt; Destinations</b> and then click the <b>Addresses</b> tab in the SNMP dialog box. |
| <b>Step 2</b> | Click <b>Create Row</b> on Fabric Manager, or click <b>Create</b> on Device Manager to add a new SNMP event destination (that is, an SNMP trap receiver).   |
| <b>Step 3</b> | Enter the IP address and port in dotted decimal notation (for example, 192.168.2.12/161) for the SNMP event destination in the <b>Address/Port</b> field.   |
| <b>Step 4</b> | Choose the SNMP protocol level from the <b>Security</b> drop-down menu.   |
| <b>Step 5</b> | Choose the SNMP event type (trap or informs). If you choose informs, set the timeout period and number of retries.  |
| <b>Step 6</b> | Click <b>Create</b> to save and apply your changes.   |
- 

## Configuring Event Security



### Caution

This is an advanced function that should only be used by administrators having experience with SNMPv3.

SNMP events can be secured against interception or eavesdropping in the same way that SNMP messages are secured. Fabric Manager or Device Manager allow you to configure the message processing model, the security model, and the security level for the SNMP events that the switch generates.

To configure SNMP event security, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | On Fabric Manager, choose <b>Switches &gt; Events &gt; SNMP Traps</b> and then click the <b>Security</b> tab in the Information pane.<br><br>On Device Manager, choose <b>Admin &gt; Events &gt; Destinations</b> on Device Manager and then click the <b>Security (Advanced)</b> tab in the SNMP dialog box. |
| <b>Step 2</b> | Set the message protocol model, security model, security name, and security level.  |
| <b>Step 3</b> | Click <b>Apply Changes</b> on Fabric Manager, or click <b>Create</b> on Device Manager to save and apply your changes.  |
-

*Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).*

## Viewing the SNMP Events Log

To view the SNMP events log from the Device Manager, choose **Logs > Events > Current** or **Logs > Events > Older**. You see the Events Log dialog box with a log of events for a single switch.

**Note**

The MDS syslog manager must be set up before you can view the event logs.

**Caution**

Changing these values from different Fabric Manager workstations at the same time may cause unpredictable results.

## Configuring RMON Using Threshold Manager

The RMON-MIB, as defined by the Internet Engineering Task Force (IETF), provides the Alarm, Log, and Events groups for monitoring appropriate statistics on the switch. The Threshold Manager in the Device Manager uses RMON in the switch to provide threshold monitoring for select statistics.

The Threshold Monitor allows you to trigger an SNMP event or log a message when the selected statistic goes over a configured threshold value. RMON calls this a rising alarm threshold. The configurable settings are:

- **Variable**—The statistic you want to set the threshold value on.
- **Value**—The value of the variable that you want the alarm to trigger at. This value is the difference (delta) between two consecutive polls of the variable by Device Manager.
- **Sample**—The sample period (in seconds) between two consecutive polls of the variable. Select your sample period such that the variable would not cross the threshold value you set under normal operating conditions.
- **Warning**—The warning level used by Device Manager to indicate the severity of the triggered alarm. This is a Fabric Manager and Device Manager enhancement to RMON.

**Note**

To configure any type of RMON alarm (absolute or delta, rising or falling threshold) click **More** on the Threshold Manager dialog box. You should be familiar with how RMON defines these concepts before configuring these advanced alarm types. Refer to the RMON-MIB (RFC 2819) for information on how to configure RMON alarms.

## Enabling RMON Alarms by Port

To configure an RMON alarm for one or more ports from the Device Manager, follow these steps:

- Step 1** Choose **Admin > Events > Threshold Manager** and click the **FC Interfaces** tab.
- Step 2** Click the **Selected** radio button to select individual ports for this threshold alarm.
  - a. Click **...** to the right of the Selected field to display all ports.
  - b. Choose the ports you want to monitor.
  - c. Click **OK** to accept the selection.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Alternatively, click the appropriate radio button to choose ports by type: **All** ports, **xE** ports, or **Fx** ports.

- Step 3** Check the check box for each variable that you want to monitor.
  - Step 4** Enter the threshold value in the Value column.
  - Step 5** Enter the sampling period in seconds. This is the time between each snapshot of the variable.
  - Step 6** Choose one of the following severity levels to assign to the alarm: Fatal, Warning, Critical, Error, Information.
  - Step 7** Click **Create**.
  - Step 8** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event. If you do not confirm the operation, the system only defines a log event.
  - Step 9** Choose **More > Alarms** from the Threshold Manager dialog box to verify the alarm you created.
- 

## Enabling RMON Alarms for VSANs

To enable an RMON alarm for one or more VSANs from the Device Manager, follow these steps:

- 
- Step 1** Choose **Admin > Events > Threshold Manager** and click the **Services** tab.  
You see the Threshold Manager dialog box with the Services tab selected.
  - Step 2** Enter one or more VSANs (multiple VSANs separated by commas) to monitor in the VSAN ID(s) field.
  - Step 3** Check the check box for each variable that you want to monitor.
  - Step 4** Enter the threshold value in the Value column.
  - Step 5** Enter the sampling period in seconds.
  - Step 6** Choose a severity level to assign to the alarm.
  - Step 7** Click **Create**.
  - Step 8** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.  
If you do not confirm the operation, the system only defines a log event.
  - Step 9** Choose **More > Alarms** from the Threshold Manager dialog box to verify the alarm you created.
- 

## Enabling RMON Alarms for Physical Components

To configure an RMON alarm for a physical component from the Device Manager, follow these steps:

- 
- Step 1** Choose **Admin > Events > Threshold Manager** and click the **Physical** tab.  
You see the Threshold Manager dialog box with the Physical tab selected.
  - Step 2** Check the check box for each variable that you want to monitor.
  - Step 3** Enter the threshold value in the Value column.
  - Step 4** Enter the sampling period in seconds.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Choose one of the following severity levels to assign to the alarm: Fatal, Warning, Critical, Error, Information.
  - Step 6** Click **Create**.
  - Step 7** Confirm the operation to define an alarm and a log event when the system prompts you to define a severity event.  
If you do not confirm the operation, the system only defines a log event.
  - Step 8** Choose **More > Alarms** from the Threshold Manager dialog box to verify the alarm you created.
- 

## Managing RMON Events

To define customized RMON events from the Device Manager, follow these steps:

- Step 1** Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.
  - Step 2** Click the **Events** tab on the RMON Thresholds dialog box.  
You see the RMON Events dialog box.
  - Step 3** Click **Create** to create a new event entry.  
You see the Create Threshold Event Entry dialog box.
  - Step 4** Configure the RMON threshold event attributes by choosing the type of event (log, snmptrap, or logandtrap).
  - Step 5** Click **Create**.
- 

## Managing RMON Alarms

To view the alarms that have already been enabled from the Device Manager, follow these steps:

- Step 1** Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.
  - Step 2** Click the **Alarms** tab.  
You see the RMON Alarms dialog box.
  - Step 3** Click **Create** to create a customized threshold entry.  
You see the Create RMON Alarms dialog box.
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Viewing the RMON Log

To view the RMON log from the Device Manager, follow these steps:

- 
- Step 1** Choose **Admin > Events > Threshold Manager** and click **More** on the Threshold Manager dialog box.
- Step 2** Click the **Log** tab on the RMON Thresholds dialog box.
- You see the RMON Log dialog box. This is the log of RMON events that have been triggered by the Threshold Manager.
-





***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



## **PART 6**

### **Network Troubleshooting**



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***



# CHAPTER 34

## Troubleshooting Your Fabric

---

There are several things you can do to use Fabric Manager to troubleshoot your fabric.

This chapter contains the following topics:

- [Troubleshooting Tools and Techniques, page 34-1](#)
- [Analyzing Switch Device Health, page 34-3](#)
- [Analyzing Switch Fabric Configuration, page 34-5](#)
- [Analyzing End-to-End Connectivity, page 34-6](#)
- [Configuring a Fabric Analyzer, page 34-7](#)
- [Using Traceroute and Other Troubleshooting Tools, page 34-13](#)
- [Analyzing the Results of Merging Zones, page 34-13](#)
- [Issuing the Show Tech Support Command, page 34-14](#)
- [Locating Other Switches, page 34-15](#)
- [Getting Oversubscription Information in Device Manager, page 34-16](#)

## Troubleshooting Tools and Techniques

Multiple techniques and tools are available to monitor and trouble shoot the Cisco MDS 9000 Family of switches. These tools provide a complete, integrated, multi-level analysis solution.

**Fabric Manager Server**—The Cisco Fabric Manager Server provides a long-term, high level view of storage network performance. Fabric wide performance trends can be analyzed using Performance Manager. It provides the starting point for deeper analysis to resolve network hot-spots.

**Device Manager**—If a performance problem is detected with the Fabric Manager Server, the Cisco Device Manager can be used to view port level statistics in real-time. Details on protocols, errors, discards, byte and frame counts are available. Samples can be taken as frequently as every 2 seconds, and values can be viewed in text form or graphically as pie, bar, area and line changes.

**Traffic Analyzer**—Another option is to launch the Cisco Traffic Analyzer for Fibre Channel from the Fabric Manager Server to analyze the traffic in greater depth. The Cisco Traffic Analyzer allows you to breakdown traffic by VSANs and protocols and to examine SCSI traffic at a logical unit number (LUN) level.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Protocol Analyzer**—If even deeper investigation is needed, the Cisco Protocol Analyzer for Fibre Channel can be launched in-context from the Cisco Traffic Analyzer. The Cisco Protocol Analyzer enables you to examine actual sequences of Fibre Channel frames easily using the Fibre Channel and SCSI decoders Cisco developed for Ethernet.

**Port Analyzer Adapter**—Fabric Manager Server and Device Manager use SNMP to gather statistics. They fully utilize the built in MDS statistics counters. Even so, there are limits to what the counters can collect.

Integration with the Cisco Traffic Analyzer and Cisco Protocol Analyzer extend the MDS analysis capabilities by analyzing the Fibre Channel traffic itself. The Cisco MDS 9000 Family Switched Port Analyzer (SPAN) enables these solutions via a flexible, non-intrusive technique to mirror traffic selectively from one or more ports to another MDS port within a fabric.

The Cisco Port Analyzer Adapter (PAA) encapsulates SPAN traffic in an Ethernet header for transport to a PC or workstation for analysis. Both Fibre Channel control and data plane traffic are available using SPAN. The PAA broadcasts the Ethernet packets, so they cannot be routed across IP networks. Hubs and switches can be used, provided they are in the same Ethernet subnet. Direct connections between a PAA and the PC are also supported. The PAA can reduce Ethernet traffic by truncating Fibre Channel data.

Both the Cisco Traffic Analyzer and Cisco Protocol Analyzer required the PAA to transport MDS SPAN traffic to a PC or workstation.



**Note**

The Cisco Traffic Analyzer works best with the Cisco Port Analyzer Adapter 2, because it provides a length value for truncated data, enabling accurate byte count reporting.

## Cisco Traffic Analyzer

The Cisco Traffic Analyzer for Fibre Channel provides real-time analysis of SPAN traffic or traffic captured previously using the Cisco Protocol Analyzer. The Fibre Channel traffic from multiple Cisco Port Analyzer Adapters (PAA) can be aggregated and analyzed by the Cisco Traffic Analyzer.

There are limits to how many SPAN sources can be sent to a single SPAN destination port on an MDS. Aggregation extends the amount of information that can be analyzed in a unified set of reports by the Cisco Traffic Analyzer.



**Note**

The aggregation capabilities are restricted to the information collect by Ethernet connections to a single PC. Aggregation across multiple PCs is NOT available.

The Cisco Traffic Analyzer presents its reports through a Web server, so you can view them locally or remotely. The traffic analysis functions are provided by 'ntop' open-source software, which was enhanced by Cisco to add Fibre Channel and SCSI analysis and MDS enhanced inter-switch link (ISL) header support for SPAN. ntop is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. ntop is also available on the Internet at <http://www.ntop.org/ntop.html>. The Cisco enhanced ntop runs under Microsoft Windows and Linux operating systems.

The Cisco Traffic Analyzer for Fibre Channel presents reports with network wide statistics. The Summary Traffic report shows what percentage of traffic was within different ranges of frames sizes. A breakdown of the percentage of traffic for each protocol like SCSI, ELS, etc. is provided. The average and peak throughput for the SPAN traffic being analyzed are also provided.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Fibre Channel traffic can be analyzed on a per VSAN basis with the Cisco Traffic Analyzer. The Domain Traffic Distribution graphs indicate how much traffic (bytes) were transmitted or received by a switch for a particular VSAN. FC Traffic Matrix graphs show how much traffic is transmitted and received between Fibre Channel sources and destinations. The total byte and frame counts for each VSAN are also provided.

Statistics can be analyzed for individual host and storage ports. You can see the percentage of SCSI read vs. write traffic, SCSI vs. other traffic, and percentage of transmitted vs. received bytes and frames. The peak and average throughput values are available for data transmitted and received by each port.

## Cisco Protocol Analyzer

The Cisco Protocol Analyzer for Fibre Channel enables you to view Fibre Channel traffic frames in real-time or from a capture file. Fibre Channel and SCSI decoders enable you to view and analyze traffic at the frame level. It matches response with request for complete decoding, which greatly simplifies navigation. Response time between response and status are presented.

The Cisco Protocol Analyzer is VSAN aware, so VSANs can be used as criteria for capture and display filters, and to colorize the display. VSAN #s can also be displayed in a column. Summary statistics are available for protocol distribution percentages and total bytes/frames transferred between specific Fibre Channel source/destination pairs. File capture and filtering controls are available. Captured files can be analyzed by either the Cisco Protocol Analyzer or the Cisco Traffic Analyzer.

Numerous features have been included for ease-of-use. You can find frames that meet particular criteria and mark them. Entries in the frame (packet) list can be colorized to highlight items of interest, and columns can be added/removed as desired.

The protocol analysis functions are provided by 'Ethereal' open-source software, which was enhanced by Cisco to decode Fibre Channel and SCSI protocols and support MDS enhanced inter-switch link (ISL) headers for SPAN. Ethereal is available on the Cisco.com software download center, under the Cisco Port Analyzer Adapter. Ethereal is also available on the Internet at <http://www.ethereal.com>. Ethereal runs under Microsoft Windows, Solaris, and Linux operating systems.

## Analyzing Switch Device Health

The Switch Health option lets you determine the status of the components of a specific switch.

To use the Switch Health option, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Select <b>Switch Health</b> from the Fabric Manager <b>Tools</b> menu.<br>You see the Switch Health Analysis window.                                      |
| <b>Step 2</b> | Click <b>Start</b> to identify any problems that may currently be affecting the selected switch.<br>You see any problems affecting the selected switches. |
| <b>Step 3</b> | Click <b>Clear</b> to remove the contents of the Switch Health Analysis window.   |
| <b>Step 4</b> | Click <b>Close</b> to close the window.   |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Online System Health Management

The Online Health Management System (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family as of Cisco MDS SAN-OS Release 1.3(4) and later.

The system health application runs on all Cisco MDS modules and runs multiple tests on each module to test individual module components and system hardware. The tests run at preconfigured intervals, cover all major fault points, and isolate any failing component in the MDS switch. The system health running on the active supervisor maintains control over all other system health components running on all other modules in the switch. The system health application running in the standby supervisor module only monitors the standby supervisor module—if that module is available in the HA standby mode.

On detecting a fault, the system health application attempts the following recovery actions:

- Sends Call Home and system messages and exception logs as soon as it detects a failure.
- Shuts down the failing module or component (such as an interface).
- Isolates failed ports from further testing.
- Reports the failure to the appropriate software component.
- Switches to the standby supervisor module if an error is detected on the active supervisor module, and a standby supervisor module exists in the Cisco MDS switch. After the switchover, the new active supervisor module restarts the active supervisor tests.
- Reloads the switch if a standby supervisor module does not exist in the switch.
- Provides CLI support to view, test, and obtain test run statistics or change the system health test configuration on the switch.
- Performs tests to focus on the problem area.
- Retrieves its configuration information from persistent storage.

Each module is configured to run the test relevant to that module. You can change the default parameters of the test in each module as required.

By default, the system health feature is enabled in each switch in the Cisco MDS 9000 Family.

## Loopback Test Configuration Frequency

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured frequency—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run at frequencies ranging from 5 seconds (default) to 255 seconds. The configured value is used for all modules. To configure the frequency of loopback tests, refer to the *Cisco MDS 9000 Family Configuration Guide*.

## Performing Internal Loopbacks

Internal loopback tests send and receive FC2 frames to and from the same ports and provides the round trip time taken in microseconds. These tests are available for both Fibre Channel and iSCSI interfaces.

Choose **Interface > Diagnostics > Internal** to perform an internal loopback test from Device Manager.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Performing External Loopbacks

External loopback tests send and receive FC2 frames to and from the same port. You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. This test is only available for Fibre Channel interfaces.

Choose **Interface > Diagnostics > External** to perform an external loopback test from Device Manager.

## Hardware Failure Action

By default, no action is taken if a failure is determined and the failed component is isolated from further testing. Failure action is controlled at individual test levels (per module), at the module level (for all tests), or for the entire switch. To configure a failure action for the switch, refer to the *Cisco MDS 9000 Family Configuration Guide*.

## Analyzing Switch Fabric Configuration

The Fabric Configuration option lets you analyze the configuration of a switch by comparing the current configuration to a specific switch or to a policy file. You can save a switch configuration to a file and then compare all switches against the configuration in the file.

To use the Fabric Configuration option to analyze the configuration of a switch, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Click <b>Fabric Configuration</b> from the Fabric Manager <b>Tools</b> menu.<br>You see the Fabric Configuration Analysis dialog box.  |
| <b>Step 2</b> | Decide whether you want to compare the selected switch to another switch, or to a Policy File. <ul style="list-style-type: none"><li>• If you are making a switch comparison, select <b>Policy Switch</b> and then click the drop-down arrow to see a list of switches.</li><li>• If you are making a policy comparison, select <b>Policy File</b>. Then click the button to the right of this option to browse your file system and select a policy file (*.XML).</li></ul> |
| <b>Step 3</b> | Click <b>Rules...</b> to set the rules to apply when running the Fabric Configuration Analysis tool.<br>You see the Rules window.  |
| <b>Step 4</b> | Change the default rules as required and click <b>OK</b> .   |
| <b>Step 5</b> | Click <b>Compare</b> .<br>The system analyzes the configuration and displays issues that arise as a result of the comparison.  |
| <b>Step 6</b> | Click to place a checkmark in the Resolve column for the issues you want to resolve.   |
| <b>Step 7</b> | Resolve them by clicking <b>Resolve Issues</b> .   |
| <b>Step 8</b> | Click <b>Clear</b> to remove the contents of the window.   |
| <b>Step 9</b> | Click <b>Close</b> to close the window.  |
-

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Analyzing End-to-End Connectivity

You can use the End to End Connectivity option to determine connectivity and routes among devices with the switch fabric. The connectivity tool checks to see that every pair of end devices can talk to each other, using a Ping test and by determining if they are in the same VSAN or in the same active zone. This option uses versions of the ping and traceroute commands modified for Fibre Channel networks.

- End to End Connectivity

The ping and redundancy tests are now mutually exclusive, you cannot run both at the same time.

To use this option, follow these steps:

- 
- Step 1** Choose **End to End Connectivity** from the Fabric Manager **Tools** menu.  
You see the End to End Connectivity Analysis dialog box.
  - Step 2** Select the VSAN in which you want to verify connectivity from the VSAN dropdown list.
  - Step 3** Select whether you want to perform the analysis for all active zones or for the default zone.
  - Step 4** Click **Ensure that members can communicate** to perform a Fibre Channel ping between the selected end points.
  - Step 5** Identify the number of packets, the size of each packet, and the timeout in milliseconds.
  - Step 6** Analyze the redundant paths between endpoints by checking the **Ensure that redundant paths exist between members** checkbox.
  - Step 7** Check the **Report errors for** checkbox to see a report of zone and device errors.
  - Step 8** Click **Analyze**.  
The End to End Connectivity Analysis window displays the selected end points with the switch to which each is attached, and the source and target ports used to connect it.  
The output shows all the requests which have failed. The possible descriptions are:
    - Ignoring empty zone—No requests are issued for this zone.
    - Ignoring zone with single member—No requests are issued for this zone.
    - Source/Target are unknown—No nameserver entries exist for the ports or we have not discovered the port during discovery.
    - Both devices are on the same switch.
    - No paths exist between the two devices.
    - VSAN does not have an active zone set and the default zone is denied.
    - Average time ... micro secs—The latency value was more than the threshold supplied.
  - Step 9** Click **Clear** to remove the contents of the window.
  - Step 10** Click **Close** to close the window.
-



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Configuring a Fabric Analyzer

Fibre Channel protocol analyzers capture, decode, and analyze frames and ordered sets on a link. Existing Fibre Channel analyzers can capture traffic at wire rate speed. They are expensive and support limited frame decoding. Also, to snoop traffic, the existing analyzers disrupt the traffic on the link while the analyzer is inserted into the link.

Cisco has brought protocol analysis within a storage network to a new level with the Cisco Fabric Analyzer. You can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be local to the point of analysis.

The Cisco Fibre Channel protocol analyzer is based on two popular public-domain software applications:

- libpcap—See <http://www.tcpdump.org>.
- Ethereal—See <http://www.ethereal.com>.



### Note

The Cisco Fabric Analyzer is useful in capturing and decoding control traffic, not data traffic. It is suitable for control path captures, and is not intended for high-speed data path captures.

This section explains the following topics:

- [About the Cisco Fabric Analyzer, page 34-7](#)
- [Configuring the Cisco Fabric Analyzer, page 34-9](#)
- [Displaying Captured Frames, page 34-10](#)


## About the Cisco Fabric Analyzer

The Cisco Fabric Analyzer comprises of two separate components (see [Figure 34-1](#)):

- Software that runs on the Cisco MDS 9000 Family switch and supports two modes of capture:
  - A text-based analyzer that supports local capture and decodes captured frames
  - A daemon that supports remote capture
- GUI-based client that runs on a host that supports libpcap such as Windows or Linux and communicates with the remote capture daemon in a Cisco MDS 9000 Family switch.

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Figure 34-1 Cisco Fabric Analyzer Usage**



## Local Text-Based Capture

This component is a command-line driven text-based interface that captures traffic to and from the supervisor module in a Cisco MDS 9000 Family switch. It is a fully functional decoder that is useful for quick debug purposes or for use when the remote capture daemon is not enabled. Additionally, because this tool is accessed from within the Cisco MDS 9000 Family switch, it is protected by the roles-based policy that limits access in each switch.

## Remote Capture Daemon

This daemon is the server end of the remote capture component. The Ethereal analyzer running on a host is the client end. They communicate with each other using the Remote Capture Protocol (RPCAP). RPCAP uses two endpoints, a TCP-based control connection and a TCP or UDP-based data connection based on TCP (default) or UDP. The control connection is used to remotely control the captures (start or stop the capture, or specify capture filters). Remote capture can only be performed to explicitly configured hosts. This technique prevents an unauthorized machine in the network from snooping on the control traffic in the network.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

RPCAP supports two setup connection modes based on firewall restrictions.

- Passive mode (default)—The configured host initiates connection to the switch. Multiple hosts can be configured to be in passive mode and multiple hosts can be connected and receive remote captures at the same time.
- Active mode—The switch initiates the connection to a configured host—one host at a time.

Using capture filters, you can limit the amount of traffic that is actually sent to the client. Capture filters are specified at the client end—on Ethereal, not on the switch.

## GUI-Based Client

The Ethereal software runs on a host, such as a PC or workstation, and communicates with the remote capture daemon. This software is available in the public domain from <http://www.ethereal.com>. The Ethereal GUI front-end supports a rich interface such as a colorized display, graphical assists in defining filters, and specific frame searches. These features are documented on Ethereal's website.

While remote capture through Ethereal supports capturing and decoding Fibre Channel frames from a Cisco MDS 9000 Family switch, the host running Ethereal does not require a Fibre Channel connection to the switch. The remote capture daemon running on the switch sends the captured frames over the out-of-band Ethernet management port. This capability allows you to capture and decode Fibre Channel frames from your desktop or laptop.

## Configuring the Cisco Fabric Analyzer

You can configure the Cisco Fabric Analyzer to perform one of two captures.

- Local capture—The command setting to enable a local capture cannot be saved to persistent storage or synchronized to standby. Launches the textual version on the fabric analyzer directly on the console screen. The capture can also be saved on the local file system.
- Remote capture—The command setting to enable a remote capture can be saved to persistent storage. It can be synchronized to the standby supervisor module and a stateless restart can be issued, if required.

To use the Cisco Fabric Analyzer feature, traffic should be flowing to or from the supervisor module.

## Sending Captures to Remote IP Addresses



### Caution

You must use the eth2 interface to capture control traffic on a supervisor module.

To capture remote traffic, use one of the following options:

- The capture interface can be specified in Ethereal as the remote device:

```
rpcap://<ipaddress or switch hostname>/eth2
```

For example:

```
rpcap://cp-16/eth2
```

```
rpcap://17.2.1.1/eth2
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- The capture interface can be specified either in the capture dialog box or by using the `-i` option at the command line when invoking Ethereal.

```
ethereal -i rpcap://<ipaddress|hostname>[:<port>]/<interface>
```

For example:

```
ethereal -i rpcap://172.22.1.1/eth2
```

or

```
ethereal -i rpcap://customer-switch.customer.com/eth2
```



**Note**

For example, in a Windows 2000 setup, click **Start** on your desktop and select **Run**. In the resulting Run window, type the required command line option in the Open field.

## Displaying Captured Frames

You can selectively view captured frames by using the display filters feature. For example, instead of viewing all the frames from a capture, you may only want to view ELP request frames. This feature only limits the captured view—it does not affect the captured or the saved frames. Procedures to specify, use, and save display filters are already documented in the Ethereal website (<http://www.ethereal.com>). Some examples of how you can use this feature are as follows:

- To view all packets in a specified VSAN, use this expression:

```
mdshdr.vsan == 2
```

- To view all SW\_ILS frames, use this expression:

```
fcswils
```

- To view class F frames, use this expression:

```
mdshdr.sof == SOFf
```

- To view all FSPF frames, use this expression:

```
swils.opcode == HLO || swils.opcode == LSU || swils.opcode == LSA
```

- To view all FLOGI frames, use this expression:

```
fcels.opcode == FLOGI
```

- To view all FLOGI frames in VSAN 1, use this expression:

```
fcels.opcode == FLOGI && mdshdr.vsan == 2
```

- To view all name server frames, use this expression:

```
dns
```

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Defining Display Filters

Display filters limit the frames that can be displayed, but not what is captured (similar to any view command). The filters to be displayed can be defined in multiple ways in the GUI application:

- Auto-definition
- Manual definition
- Assisted manual definition
- Only manual definition in local capture
- No assists

Regardless of the definition, each filter must be saved and identified with a name.



### Note

This GUI-assisted feature is part of Ethereal and you can obtain more information from <http://www.ethereal.com>.

## Capture Filters

You can limit what frames are captured by using the capture filters feature in a remote capture. This feature limits the frames that are captured and sent from the remote switch to the host. For example, you can capture only class F frames. Capture filters are useful in restricting the amount of bandwidth consumed by the remote capture.

Unlike display filters, capture filters restrict a capture to the specified frames. No other frames are visible until you specify a completely new capture.

The syntax for capture filter is different from the syntax for display filters. Capture filters use the Berkeley Packet Filter (BPF) library that is used in conjunction with the libpcap freeware. The list of all valid Fibre Channel capture filter fields are provided later in this section.

Procedures to configure capture filters are already documented in the Ethereal website (<http://www.ethereal.com>). Some examples of how you can use this feature as follows:

- To capture frames only on a specified VSAN, use this expression:  
`vsan = 1`
- To capture only class F frames, use this expression:  
`class_f`
- To capture only class Fibre Channel ELS frames, use this expression:  
`els`
- To capture only name server frames, use this expression:  
`dns`
- To capture only SCSI command frames, use this expression:  
`fcpcmd`



### Note

This feature is part of libpcap and you can obtain more information from <http://www.tcpdump.org>.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Permitted Capture Filters

This section lists the permitted capture filters.

- o vsan
- o src\_port\_idx
- o dst\_port\_idx
- o sof
- o r\_ctl
- o d\_id
- o s\_id
- o type
- o seq\_id
- o seq\_cnt
- o ox\_id
- o rx\_id
- o els
- o swils
- o fcp\_cmd (FCP Command frames only)
- o fcp\_data (FCP data frames only)
- o fcp\_rsp (FCP response frames only)
- o class\_f
- o bad\_fc
- o els\_cmd
- o swils\_cmd
- o fcp\_lun
- o fcp\_task\_mgmt
- o fcp\_scsi\_cmd
- o fcp\_status
- o gs\_type (Generic Services type)
- o gs\_subtype (Generic Services subtype)
- o gs\_cmd
- o gs\_reason
- o gs\_reason\_expl
- o dns (name server)
- o udns (unzoned name server)
- o fcs (fabric configuration server)
- o zs (zone server)
- o fc (use as fc[x:y] where x is offset and y is length to compare)
- o els (use as els[x:y] similar to fc)
- o swils (use as swils[x:y] similar to fc)
- o fcp (use as fcp[x:y] similar to fc)
- o fcct (use as fcct[x:y] similar to fc)

## Using the Ping Tool

You can use the Ping tool to determine connectivity from another switch to a port on your switch.

To use the Ping tool, follow these steps:

- 
- Step 1** Select **Ping** from the Fabric Manager **Tools** menu. You can also select it from the right-click context menus for hosts and storage devices in the Fabric pane.  
You see the Ping dialog box.
  - Step 2** Select the source switch from the Source Switch drop-down list.
  - Step 3** Select the VSAN in which you want to verify connectivity from the VSAN drop-down list.
  - Step 4** Select the target end port for which you want to verify connectivity from the Target Endport drop-down list.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Click **Start** to perform the ping between your switch and the selected port.
- Step 6** In a few seconds, you see the results in the Results area of the dialog box.
- Step 7** Click **Clear** to clear the contents of the window and perform another ping, or click **Close** to close the window.
- 

## Using Traceroute and Other Troubleshooting Tools

You can use the following options on the Tools menu to verify connectivity to a selected object or to open other management tools:

- Traceroute—Verify connectivity between two end devices that are currently selected on the Fabric pane.
- Device Manager—Launch the Device Manager for the switch selected on the Fabric pane.
- Command Line Interface—Open a Telnet or SSH session for the switch selected on the Fabric pane.

To use the Traceroute option to verify connectivity, follow these steps:

- 
- Step 1** Select **Traceroute** from the Fabric Manager **Tools** menu.
- You see the Traceroute dialog box.
- Step 2** Select the source switch from the Source Switch drop-down list.
- Step 3** Select the VSAN in which you want to verify connectivity from the VSAN drop-down list.
- Step 4** Select the target end port for which you want to verify connectivity from the Target Endport drop-down list.
- Step 5** Click **Start** to perform the traceroute between your switch and the selected port.
- Step 6** In a few seconds, you see the results in the Results area of the dialog box.
- Step 7** Click **Clear** to clear the contents of the window and perform another traceroute, or click **Close** to close the window.
- 

## Analyzing the Results of Merging Zones

You can use the Zone Merge option on the Fabric Manager Zone menu to determine if two connected switches have compatible zone configurations.

To use the Zone Merge option, follow these steps:

- 
- Step 1** Choose **Merge Analysis** from the Fabric Manager **Zone** menu.
- The Zone Merge Analysis dialog is displayed.
- Step 2** Select a switch from each drop-down list.
- Step 3** Select the VSAN for which you want to perform the zone merge analysis.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 4** Click **Analyze**.

The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches.

**Step 5** Click **Clear** to remove the contents of the window.

**Step 6** Click **Close** to close the window.

---

## Issuing the Show Tech Support Command

The **show tech support** command is useful when collecting a large amount of information about your switch for troubleshooting purposes. The output can be provided to technical support representatives when reporting a problem.

You can issue a **show tech support** command from Fabric Manager for one or more switches in a fabric. The results of each command are written to a text file, one file per switch, in a directory you specify. You can then view these files using Fabric Manager.

You can also save the Fabric Manager map as a JPG file. The file is saved with the name of the seed switch (for example, 172.22.94.250.jpg).

You can zip up all the files (the show tech support output and the map file image) and send the resulting zipped file to technical support.

To issue the **show tech support** command in Fabric Manager, follow these steps.

---

**Step 1** Select **Show Tech Support** from the **Tools** menu.

You see the Show Tech Support dialog box.

**Step 2** Select the switches for which you want to view Show Tech Support information by checking the check boxes next to their IP addresses.

**Step 3** Set the timeout value.

The default is 30 seconds.

**Step 4** Select the folder where you want the text files (containing the Show Tech Support information) to be written.

**Step 5** Check the **Save Map** check box if you want to save a screenshot of your map as a JPG file.

**Step 6** Check the **Compress all files as** check box to compress the files into a zip file.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Step 7** Click the **OK** button to start issuing the show tech support command to the switches you specified, or click the **Close** button to close the Show Tech Support dialog box without issuing the show tech support command.

In the Status column next to each switch, a highlighted status is displayed. A yellow highlight indicates that the Show Tech Support command is currently running on that switch. A red highlight indicates an error. A green highlight indicates that the Show Tech Support command has completed successfully. On successful completion, a button becomes available in the View column for each switch.

**Step 8** If prompted, enter your username and password in the appropriate fields for the switch in question.

Note that in order for Fabric Manager to successfully issue the show tech support command on a switch, that switch must have this username and password. Fabric Manager is unable to log into a switch that does not have this username and password, and an error is returned for that switch.



**Note** If you would like to view the Show Tech Support files without using Fabric Manager, you can open them with any text editor. Each file is named with the switch's IP address and has a .TXT extension (for example, 111.22.33.444.txt).

## Locating Other Switches

The Locate Switches option uses SNMPv2 and discovers devices responding to SNMP requests with the read-only community string public. You can use this feature if:

- You have third-party switches that do not implement the FC-GS3 FCS standard that provides management IP addresses.
- You want to locate other Cisco MDS 9000 switches in the subnet but are not physically connected to the fabric (and therefore cannot be found via neighbors).

To locate switches that are not included in the currently discovered fabric, follow these steps:

**Step 1** Select **Locate Switches and Devices** from the Fabric Manager **File** menu.

You see the Locate Switches dialog box.

**Step 2** In the Comma Separated Subnets field, enter a range of specific addresses belonging to a specific subnet which limit the research for the switches. To look for a Cisco MDS 9000 switch belonging to subnet 192.168.199.0, use the following string:

**192.168.100.[1-254]**

Multiple ranges can be specified, separated by commas. For example, to look for all the devices in the two subnets 192.168.199.0 and 192.169.100.0, use the following string:

**192.168.100.[1-254], 192.169.100.[1-254]**

**Step 3** Enter the appropriate read community string in the Read Community field.

The default value for this string is "public."

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 4** Click **Display Cisco MDS 9000 Only** to display only the Cisco MDS 9000 Family switches in your network fabric.
- Step 5** Click **Search** to discover switches and devices in your network fabric. You see the results of the discovery in the Locate Switches window.

**Note**

The number in the lower left corner of the screen increments as the device locator attempts to discover the devices in your network fabric. When the discovery process is complete, the number indicates the number of rows displayed.

## Getting Oversubscription Information in Device Manager

To determine the oversubscription for a module using Device Manager, follow these steps:

- Step 1** Right-click the module you want to check oversubscription on and select **Check Oversubscription** from the pop-up menu. You see the oversubscription dialog box.
- Step 2** Click **Close** to close this dialog box.



## CHAPTER 35

# Management Software Troubleshooting

This chapter answers some of the most frequently asked questions about Cisco Fabric Manager and Device Manager. This chapter contains the following topics:

- **Installation Issues**
  - When installing Fabric Manager from windows, clicking the install button fails., page 35-3
  - How do I install Java Web Start on a UNIX machine?, page 35-4
  - Why can't I launch Fabric Manager on Solaris?, page 35-4
  - Why is my browser prompting to save JNLP files?, page 35-4
  - Why do I get a "Java Web Start not detected" error?, page 35-4
  - Why can't I see my desktop shortcuts?, page 35-5
  - How do I upgrade to a newer version?, page 35-5
  - How do I downgrade Fabric Manager or Device Manager?, page 35-5
  - What do I do if my upgrade is not working?, page 35-5
  - Java Web Start hangs on download dialog. What do I do?, page 35-5
  - How can I manually configure my browser for Java Web Start?, page 35-6
  - Can I run Java Web Start from the command line?, page 35-6
  - Windows 2000 crashes (blue screen). What do I do?, page 35-6
  - How do I clear the Java Web Start cache?, page 35-7
  - Why doesn't my login work in Fabric Manager and Device Manager?, page 35-7
  - Why can't I install Fabric Manager or Device Manager when pcAnywhere is running?, page 35-7
  - The Fabric Manager or the Performance Manager service shows up as "disabled" in the Services menu., page 35-7
  - Why can't I install Fabric Manager or Device Manager when McAfee Internet Suite 6.0 Professional is running?, page 35-8
  - I get an error ".sm/logon." when I downgrade from MDS SAN-OS Release 2.x (or newer) to 1.3(x)., page 35-8
- **General**
  - Why do I get errors while monitoring Area chart graphing?, page 35-8
  - Why do I get "gen error" messages?, page 35-8

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Why are disk images in the Device Manager Summary View not showing up?, page 35-8
- Why can't I set both the D\_S\_TOV and E\_D\_TOV timers in the Device Manager?, page 35-9
- Why are the columns in the Device Manager tables too small?, page 35-9
- Why are my fabric changes not propagated onto the map (for example, links don't disappear)?, page 35-9
- Why does the PortChannel creation dialog become too small after several uses?, page 35-9
- Why do I see errors when I have configured IPFC?, page 35-9
- Why is Fabric Manager or Device Manager using the wrong network interface?, page 35-9
- Why am I seeing display anomalies?, page 35-10
- How do I connect the Fabric Manager client to the server across VPN?, page 35-10
- Why is the active zone set in edit zone always shown in bold (even after successful activation)?, page 35-10
- Can I create a zone with prefix IVRZ and a zone set with name nozonset?, page 35-10
- One-click license install fails, cannot connect to Cisco website., page 35-10
- Fabric Manager client and Device Manager cannot connect to the switch, page 35-10
- License Wizard fails to fetch license keys, saying connect failed, page 35-11
- Windows Issues
  - Text fields showing up too small, cannot enter any data, page 35-11
  - Why does CiscoWorks fail to start in the browser?, page 35-11
  - Help contents are unreadable because of highlighting, page 35-11
  - Printing causes an application crash, page 35-11
  - Windows XP hangs (or blue screen). What do I do?, page 35-12
  - Why do the Device Manager Icons Disappear Sometimes?, page 35-12
  - Why does Fabric Manager hang when I drag an existing Zone Member to a Zone?, page 35-12
  - Why does SCP/SFTP fail when I try to copy a file from my local machine to the switch?, page 35-12
- UNIX Issues
  - Why Do the Parent Menus Disappear?, page 35-13
  - Why do I keep getting a "too many open files" error?, page 35-13
- Other
  - How can I set the map layout so it stays after I restart Fabric Manager?, page 35-14
  - Two switches show on my map, but I only have one switch, page 35-14
  - There is a red/orange/dotted line through the switch. What's wrong?, page 35-14
  - Can I upgrade without losing my map settings?, page 35-19
  - I see "Please insure that FM server is running on localhost.", page 35-20
  - How can I run Cisco Fabric Manager if I have multiple interfaces?, page 35-21
  - How can I configure an HTTP proxy server?, page 35-22
  - How can I clear the topology map?, page 35-23

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- [Can I use Fabric Manager in a mixed software environment?, page 35-23](#)
- [I Get an Error When Launching Fabric Manager, page 35-23](#)
- [Can I Search for Devices in a Fabric?, page 35-24](#)
- [Do I Need A License of Fabric Manager Server for Each Switch in the Fabric?, page 35-24](#)
- [How can I Manage Multiple Fabrics?, page 35-24](#)
- [License Expiration Causes Orange X Through Switch, page 35-24](#)

## Installation Issues

### When installing Fabric Manager from windows, clicking the install button fails.

Make sure that Java Web Start is installed properly. To check, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Go to the Programs menu and see if Java Web Start is there.  |
| <b>Step 2</b> | Start the <b>Java Web Start</b> program to make sure there is no problem with the Java Runtime installation.   |
| <b>Step 3</b> | Click the <b>Preferences</b> tab, and make sure the proxies settings are fine for Web Start.   |
| <b>Step 4</b> | Check that your browser is set up to handle jnlp settings properly (see the <a href="#">“How can I manually configure my browser for Java Web Start?”</a> section on page 35-6). |
- 

If you had older versions of the application and you see an error pop-up window saying cannot open the JNLP file (in the error details), this could be because the Java Web Start cache is messed up. To work around this, clear the cache and retry. To clear the cache, see the [“How do I clear the Java Web Start cache?”](#) section on page 35-7.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## How do I install Java Web Start on a UNIX machine?

If you are using UNIX (Linux, Solaris) the Sun JRE 1.4.0 and 1.4.1 does not automatically install Java Web Start. However, the Web Start zip file is bundled with the JRE.

To install Java Web Start, follow these steps:

---

**Step 1** In the directory where you have installed the JRE, there is a Java Web Start zip file (it is named something like **javaws-1\_2\_linux-i586-i.zip**).

**Step 2** Unzip this file and run the install.sh script.

You are prompted to enter the path to the java installation.

**Step 3** Update the mime-type settings for users. This can be done for all users.

```
# /etc/mime.types should contain the line
    type=application/x-java-jnlp-file desc="Java Web Start" exts="jnlp"
```

```
# /etc/mailcap should contain the line
    application/x-java-jnlp-file; /javaws %s
```

To install for individual users add the lines to \$HOME/.mime.types and \$HOME/mailcap.

---

## Why can't I launch Fabric Manager on Solaris?

If you are using Solaris 2.8 and are logged in as root and are using Netscape Navigator 6, you will not be able to register the mime-type. Regular users can register the mime-type with Netscape Navigator 6 by manually adding it. Netscape 4.x works fine for all users.

## Why is my browser prompting to save JNLP files?

Your browser may not be set up to launch Java Web Start for JNLP mime types. Java Web Start is probably not installed or configured properly (see the [“How can I manually configure my browser for Java Web Start?”](#) section on page 35-6).

## Why do I get a “Java Web Start not detected” error?

If you installed Java Web Start but still see an error message (in red) saying "Java Web Start not detected..." on the switch home page, it could be a simple JavaScript error. We try to detect a Java Web Start installation by running some JavaScript code tested for Internet Explorer and Mozilla (newer versions). On some browsers (for example, Netscape 6.0, Opera) this code does not work properly although the links still work.

- First, try clicking on the install links.
- If that does not work, check to see if the browser helper applications settings are correct (for example, for Netscape 6.0 **Edit > Preferences > Navigator > Helper Applications**). See the [“How can I manually configure my browser for Java Web Start?”](#) section on page 35-6.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Why can't I see my desktop shortcuts?

For Windows 2000 and Windows NT, we create Program Menu entries (under a new Cisco MDS 9000 program menu) and desktop shortcuts for Fabric Manager and Device Manager. The desktop shortcuts and start menu entries for Fabric Manager and Device Manager are called FabricManager and DeviceManager respectively. In other versions of Windows, including XP, we just create batch files on the desktop called FabricManager.bat and DeviceManager.bat. For UNIX, we create shell scripts called FabricManager.sh and DeviceManager.sh under the \$HOME/.cisco\_mds9000/bin directory. Note that on Windows, installations run under Mozilla variants of browsers, and the desktop shortcuts do not get created. The workaround is to manually create desktop shortcuts.

## How do I upgrade to a newer version?

To upgrade to a newer version of Fabric Manager or Device Manager, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Be sure all running instances of Fabric Manager or Device Manager are closed.   |
| <b>Step 2</b> | Point your browser at the switch running the new version and click on the appropriate install link. Fabric Manager or Device Manager prompts you to upgrade if the switch is running a newer version. |
- The installer checks your local copies and updates any newer versions of the software.
- 

## How do I downgrade Fabric Manager or Device Manager?

Cisco MDS SAN-OS Release 2.x or later supports downgrades using the installer. For earlier releases, downgrades are not supported through the installer. To downgrade Fabric Manager or Device Manager for an earlier release, you need to manually uninstall them and then install the previous version of Fabric Manager or Device Manager. See the [“Downgrading the Management Software” section on page 1-9](#).

## What do I do if my upgrade is not working?

If you are trying to upgrade because Fabric Manager or Device Manager prompted you saying that the switch version is higher, and the upgrade failed, it might be because your default browser settings are incorrect. Some error must have occurred during your last browser upgrade/install. To work around this, launch the browser independently and click on install.

On rare occasions, we have seen the upgrade happen but the version does not change. This is because of HTTP caching in the network. During the upgrade, HTTP requests for files on the switch get cached in the local machine. Even though the switch is in a higher version, the management software installed is at the old version. The workaround for this is to uninstall the Fabric/Device Manager, clear the Java Web Start cache, and then do a clean install.

## Java Web Start hangs on download dialog. What do I do?

To make sure Java Web Start is set up to access the switch in the same way your browser is set up, follow these steps:

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- 
- Step 1** Start up Java Web Start (javaws.exe/javaws). You see the Java Web Start Application Manager.
- Step 2** Choose **File > Preferences > General** and make sure your proxy settings are correct. For example, if you are using an HTTP proxy, set it up here.
- Step 3** Check **Use Browser**.
- 

## How can I manually configure my browser for Java Web Start?

For browsers like Opera, certain versions of Mozilla, or Konqueror, you must manually register Java Web Start as the helper application for the JNLP files. To do this, the data you need is:

- Description=Java Web Start
- File Extension=jnlp
- Mime Type=application/x-java-jnlp-file
- Application=path-to-javaws (e.g. /usr/local/javaws/javaws)

After setting this up, you may need to restart the browser. If you see "Java Web Start not detected" warnings, you can ignore them. These warnings are based on JavaScript, and not all browsers behave well with JavaScript. Click on the install links to install Fabric Manager or Device Manager.



### Note

For Windows Users: To set up Java Web Start on \*.jnlp files, select **Windows Explorer > Tools > Folder Options > File Types**. Either change the existing setting for JNLP or add one so that \*.jnlp files are opened by javaws.exe. This executable is under Program Files\Java Web Start

## Can I run Java Web Start from the command line?

If you cannot get your browser to run Java Web Start, you can still run Java Web Start from the command line (javaws.exe or javaws) giving it the URL of the Fabric Manager or Device Manager on the switch as an argument. For example, if your switch IP address is 10.0.0.1, you would use these commands to start Fabric Manager and Device Manager:

```
javaws http://10.0.0.1/cgi-bin/fabric-manager.jnlp
javaws http://10.0.0.1/cgi-bin/element-manager.jnlp
```

## Windows 2000 crashes (blue screen). What do I do?

Be sure you have Service Pack 3 installed if you are using JRE 1.4.1. (You should actually have been prompted to install Service Pack 3 during the JRE 1.4.1 installation.) If you do not have it installed, Windows 2000 may crash. If you do not want to upgrade to Service Pack 3, you can install JRE 1.4.0.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## How do I clear the Java Web Start cache?

To clear the Java Web Start cache, follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Start the Java Web Start Application Manager (javaws.exe, javaws).  |
| <b>Step 2</b> | Go to <b>File &gt; Preferences &gt; Advanced</b> and clear the applications folder/cache. You can manually delete the .javaws/cache directory. On Windows this is under Documents and Settings, and on UNIX this is under \$HOME. |
- 

## Why doesn't my login work in Fabric Manager and Device Manager?

Make sure you have done the Initial Setup Routine on the switch. Refer to the *Cisco MDS 9000 Family Configuration Guide*. Quick checks:

- Make sure that the management interface on the switch is up (**show interface mgmt0**).
- Check whether you can connect to the management interface (**ping**).
- Verify the username is valid (**show snmp user**). You can also add/edit the users through the CLI.
- If you have multiple network interfaces, see the “[Why is Fabric Manager or Device Manager using the wrong network interface?](#)” section on page 35-9

## Why can't I install Fabric Manager or Device Manager when pcAnyWhere is running?

You can either stop the pcAnyWhere service and install Fabric Manager or Device Manager, or install/update DirectX. For more information, refer to the website at <http://forum.java.sun.com>.

## The Fabric Manager or the Performance Manager service shows up as “disabled” in the Services menu.

This could happen if:

- The service menu for Fabric Manager or Performance Manager was open during an uninstall/upgrade.
- The Fabric Manager client or Device Manager was running while doing an uninstall/upgrade.

This error happens when Windows is unable to delete a service completely. A reboot of the host should fix the problem.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Why can't I install Fabric Manager or Device Manager when McAfee Internet Suite 6.0 Professional is running?

The McAfee internet suite comes with a virus scanner, firewall, antispyware, and privacy management. The privacy management can interfere with the Fabric Manager server-client interactions. To work around this you must shut down the privacy service.

## I get an error ".sm/login." when I downgrade from MDS SAN-OS Release 2.x (or newer) to 1.3(x).

The installer does not support a downgrade from Cisco MDS SAN-OS Release 2.x (or newer) to Cisco MDS SAN-OS Release 1.3(x) or earlier. Fabric Manager and Device Manager are backwardly compatible so we suggest running the newer versions. If you still want to downgrade to the lower version, see the [“Downgrading the Management Software” section on page 1-9](#).

## General

### Why do I get errors while monitoring Area chart graphing?

When doing the area chart graphing from the monitor window, if you move the mouse over the Area chart before the first data comes back, you see a java.lang.ArrayIndexOutOfBoundsException error on the message log from JChart getX(). This is because JChart tries to locate a value that does not exist yet. This might be fixed in a future version of JChart.

### Why do I get “gen error” messages?

Usually a “gen error” means that the SNMP agent on the switch had an unexpected error in the process of serving an SNMP request. However, when you are accessing the switch through a VPN connection or any sort of NAT scheme, all errors are reported as gen error. This is a known problem and will be fixed in a future release. You can verify whether this was the reason behind your gen error by trying to reproduce this error in an environment where there is no network address translation (where you are on the same network as the switch).

### Why are disk images in the Device Manager Summary View not showing up?

On some occasions the Summary View table in the Device Manager does not show the icons for disks attached to a Fx port. This is because the FC4 features are empty for this port. A LUN discovery must be issued to discover information about these hosts/disks that do not register their FC4 types. You can do this in the Device Manager by clicking **FC > Advanced > LUNs**.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Why can't I set both the D\_S\_TOV and E\_D\_TOV timers in the Device Manager?

If you modify both E\_D\_TOV and D\_S\_TOV at the same time, and the new D\_S\_TOV value is larger than the old E\_D\_TOV value, you will get a WrongValue error. To work around this, you must change the values separately.

## Why are the columns in the Device Manager tables too small?

If Device Manager is trying to display a large table and your switch is running slowly, the table will come up with the tabs being hidden. To work around this, you must resize the window to see the data.

## Why are my fabric changes not propagated onto the map (for example, links don't disappear)?

Fabric Manager shows that a device or port is down by displaying a red cross on that port or device. However, Fabric Manager does not remove any information that's already discovered. You must rediscover to correctly update the map.

## Why does the PortChannel creation dialog become too small after several uses?

After several uses, the MemberList TextBox (in the PortChannel Create Window) does not display as it should. It changes from a long TextBox with a ComboBox for choosing ports, to a small square TextBox that is too small to choose ports. This is a known problem and will be fixed in a future release. To work around this problem, stop and restart Fabric Manager or Device Manager.

## Why do I see errors when I have configured IPFC?

When IPFC and out of band management are configured, the Device Manager might not work using SNMPv3 if you use the IPFC address. The workaround is either to use the management interface (mgmt0) address, or to use SNMPv1/v2c over IPFC.

## Why is Fabric Manager or Device Manager using the wrong network interface?

The problem happens because the underlying Java library picks a local interface arbitrarily. To work around this, you can supply a command line argument before starting the Fabric/Device Manager. In the desktop shortcut or shell script or batch file, add the following parameter "-Device Managerds.nmsAddress="

For example, in Windows the line looks like ".javaw.exe -Device Managerds.nmsAddress=X.X.X.X -cp .".

In desktop shortcuts, this length could exceed the maximum characters allowed. If this happens, delete the "-Dsun.java2d.ddoffscreen=false" portion to make more space. Newer versions of Fabric Manager (Release 1.2 and later) allow you to pick a preferred network interface.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Why am I seeing display anomalies?

If you see Fabric Manager or Device Manager submenus detached from menus, the mouse pointer in Fabric Manager Map is slow to react to mouse movement, or a wrong tooltip is displayed, these are display anomalies, not problems with Fabric Manager or Device Manager.

Some older video cards exhibit these display anomalies. To fix this, first try updating the video drivers. If this doesn't solve the problem, replace the video card.

## How do I connect the Fabric Manager client to the server across VPN?

To make the Fabric Manager client work across VPN, you must set up a Java command line option `-Djava.rmi.server.hostname=client VPN IP address`.

## Why is the active zone set in edit zone always shown in bold (even after successful activation)?

A member of this VSAN must be participating in IVR zoning. Because the IVR zones get added to active zones, the active zone set configuration is always different from the local zone set configuration with the same name. The zone set name is always bold.

## Can I create a zone with prefix IVRZ and a zone set with name nozonset?

Do not use these special names. These names are used by the system for identifying IVR zones.

## One-click license install fails, cannot connect to Cisco website.

The one-click license install tries to open an HTTP connection to the Cisco website. If you do your browsing using an HTTP proxy then the following command-line variables need to be added to your Fabric Manager client scripts:

`-Dhttps.proxyHost` and `-Dhttps.proxyPort`.

In case your one-click install URL starts with "http://" (and not "https://"), the variables are:

`-Dhttp.proxyHost` and `-Dhttp.proxyPort`.

For example, in Windows, edit the MDS 9000\bin\FabricManager.bat file and add to the JVMARGS `"-Dhttps.proxyHost=HOSTADDRESS -Dhttps.proxyPort=HOSTPORT"`.

## Fabric Manager client and Device Manager cannot connect to the switch

Fabric Manager or Device Manager using SNMPv3 at Cisco MDS SAN-OS Release 1.3(3) or earlier can't manage a switch running Release 1.3(4) or later. This might affect a software upgrade using Fabric Manager from Release 1.3(3) to Release 1.3(4).

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## License Wizard fails to fetch license keys, saying connect failed

Java versions 1.4.2\_01 and older don't seem to have the right set of CA (certifying authority) certificates to validate the SSL certificates on the EMC server (https). The license wizard is unable to make an https connection to the EMC servers. The workaround is to install the latest 1.4(x) version of Java, preferably 1.4.2\_04 or later.

## How do I increase log window size in Fabric Manager Client?

To limit the memory usage by FM Client, the log window is limited to 500 lines by default. If you want to increase this, edit `sm.properties` in `<install directory>/db/<user>` directory and change `LogBufferSize`.

## Windows Issues

### Text fields showing up too small, cannot enter any data

When Reflection X is running, certain text fields in the Fabric Manager and Device Manager are not rendered to the full width of the field. Resize the dialog box to see the text fields properly.

### Why does CiscoWorks fail to start in the browser?

CiscoWorks fails to come up in the browser. This could be because CiscoWorks does not support Java JVM 1.4.0. To turn off Java JVM 1.4.0 in Internet Explorer, select the **Tools/Internet Options** menu item, click on the **Advanced** tab, and uncheck the **Use Java 2 1.4.0** option.

### Help contents are unreadable because of highlighting

With the Windows look and feel, the highlight background color is dark blue and so the black text is unreadable. This is a known problem with Java Help. For more details refer to:  
<http://developer.java.sun.com/developer/bugParade/bugs/4727419.html>.

### Printing causes an application crash

On Windows NT there is a known Sun JVM bug - the printservice crashes the VM. The solution suggested by Sun is to update NT with SP 6. For more details refer to:  
<http://developer.java.sun.com/developer/bugParade/bugs/4530428.html>.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Windows XP hangs (or blue screen). What do I do?

Windows XP with the ATI Radeon AGP graphics cards has known to freeze (hang) when a Java application exits. The newer drivers from ATI seem to have fixed this problem. The other workaround is to run the application with "-Dsun.java2d.noddraw=true". We do this today in the shortcut and shell scripts we create. For more details refer to:

<http://developer.java.sun.com/developer/bugParade/bugs/4713003.html>.

## Why do the Device Manager icons disappear sometimes?

On certain versions of Windows, certain images disappear. This is a Java bug. We have a workaround that is already in place (disable DirectDraw acceleration) - but there are still cases where this problem might arise. For more details refer to:

<http://developer.java.sun.com/developer/bugParade/bugs/4664818.html>.

## Why does Fabric Manager hang when I drag an existing zone member to a zone?

When dragging a zone member to a zone (where that member is already present) you get an error message saying the zone member is already present and the application freezes. This is a Sun Java bug, and the problem is seen with JRE versions earlier than 1.4.2. For more details refer to

<http://developer.java.sun.com/developer/bugParade/bugs/4633417.html>. Use a Sun JRE with version 1.4.2 or later where this problem does not occur.

## Device Manager or Fabric Manager window content disappears in Windows XP

Device Manager or Fabric Manager main window content disappears in Windows XP due to a Java bug. Refer to the following website:

[http://bugs.sun.com/bugdatabase/view\\_bug.do?bug\\_id=4919780](http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=4919780).

Minimize or maximize the window and restore to the normal size to restore the window content. Disabling Direct Draw may also prevent this from happening by adding "-Dsun.java2d.noddraw=true" to JVMARGS in <FM-install-dir>/bin/FabricManager.bat and DeviceManager.bat

## Why does SCP/SFTP fail when I try to copy a file from my local machine to the switch?

If there are embedded spaces in the file path, then windows scp/sftp might fail. You will get a copyDeviceBusy error from the switch. In tools such as the License Wizard either make sure tftp copy can be done or pick filenames with no spaces.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## UNIX Issues

### Why do the parent menus disappear?

Displaying a submenu may occasionally cause the parent menu to disappear. For more details on this bug, refer to: <http://developer.java.sun.com/developer/bugParade/bugs/4470374.html>.

### Why do I keep getting a "too many open files" error?

If you are running the JVM (Java Virtual Machine) on Linux and the drive where Java is installed or your home directory is NFS mounted, there is an open bug against the Sun JDK about errors acquiring file locks. The symptoms for the Fabric Manager are that launching a Device Manager or saving/opening files will fail, giving a "too many open files" I/O or socket exception. The JVM keeps trying to open a file on the NFS mounted drives, fails, and keeps trying to do it until it hits the 1024 file descriptors limit. Workarounds (assuming /tmp is a local disk - replace it with your tmp area):

- System Preferences

Make sure the system level preferences are stored on a local disk. The system preferences are stored in \$JAVA\_HOME/.systemPrefs where JAVA\_HOME is where you have installed the JDK. If this directory is NFS mounted, then just do the following:

```
$ rm -rf $JAVA_HOME/.systemPrefs<
$ mkdir /tmp/.systemPrefs
$ ln -s /tmp/.systemPrefs $JAVA_HOME/.systemPrefs
```

The problem with this workaround is that you have to make sure /tmp/.systemPrefs exists on every box where you are using \$JAVA\_HOME. We recommend installing the JVM as root and on a local disk.

- User Preferences

If your home directory is NFS mounted and you are getting this problem. Do the following:

```
$ rm -rf $HOME/.java
$ mkdir /tmp/.java.$USER
$ ln -s /tmp/.java.$USER $HOME/.java
```

For further details, see the following URLs:

<http://developer.java.sun.com/developer/bugParade/bugs/4673298.html>

<http://developer.java.sun.com/developer/bugParade/bugs/4635353.html>

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Other

### How can I set the map layout so it stays after I restart Fabric Manager?

If you have arranged the map to your liking and would like to “freeze” the map so that the objects stay as they are even after you stop Fabric Manager and restart it again, follow these steps:

- 
- Step 1** Right-click on a blank space in the map. A menu is displayed.
- Step 2** Select **Layout > Fix All Nodes** from the menu.
- 

### Two switches show on my map, but I only have one switch

If two switches show on your map, but you only have one switch, it may be that you have two switches in a non-contiguous VSAN that have the same Domain ID. Fabric Manager uses <vsanId><domainId> to look up a switch, and this can cause the fabric discovery to assign links incorrectly between these errant switches.

The workaround is to verify that all switches use unique domain IDs within the same VSAN in a physically connected fabric. (The fabric configuration checker will do this task.)

### There is a red/orange/dotted line through the switch. What’s wrong?

If a red line shows through your switch, this means Fabric Manager sees something wrong with the switch. Choose **Switches** in the Physical Attributes pane to see a status report in the information pane. A module, fan, or power supply has failed or is offline and plugged in.

If a dotted orange line shows through your switch, this indicates a minor status warning for that switch. Usually it means an issue with one of the modules. The tooltip should say exactly what is wrong. Hold the mouse over the switch to see the tooltip.

Below are tables of color settings and tooltip definitions for Fabric Manager and Device Manager.

**Table 35-1** *Fabric Manager and Device Manager Color Definitions*

<b>Fabric Manager Color</b>	<b>Definition</b>
Red Slash	Cannot communicate with a switch via SNMP.
Red X	Cannot communicate with or see a switch in the Domain Manager/Fabric Configuration Server list of fabric switches.
<b>Device Manager Color</b>	<b>Definition</b>
Green Square with Mode (e.g., F, T, TE, U/I for FICON)	Port up.
Orange Square with Mode	Trunk incomplete.
Orange Cross	Ols or Nos received.
Brown Square	Port is administratively down.



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 35-1 Fabric Manager and Device Manager Color Definitions (continued)**

Fabric Manager Color	Definition
Light Gray Square	Port is not manageable.
Red Cross	HardwareFailure/LoopbackDiagFailure/LinkFailure
Red Square	Any other kind of configuration failure.
No Square or Black Square	Port not yet configured.

**Table 35-2 Device Manager Tooltip Definitions**

Tooltip	Definition
adminDown	The port is administratively down.
bitErrRTThresExceeded	Bit error rate too high.
bundleMisCfg	Misconfiguration in PortChannel membership detected.
channelAdminDown	This port is a member of a PortChannel and that PortChannel is administratively down.
channelConfigurationInProgress	This port is undergoing a PortChannel configuration.
channelOperSuspended	This port is a member of a PortChannel and its operational parameters are incompatible with the PortChannel parameters.
deniedDueToPortBinding	Suspended due to port binding.
domainAddrAssignFailureIsolation	The elected principal switch is not capable of performing domain address manager functions so no Nx_port traffic can be forwarded across switches, hence all Interconnect_Ports in the switch are isolated.
domainInvalidRCFReceived	Invalid RCF received.
domainManagerDisabled	Domain manager is disabled.
domainMaxReTxFailure	Domain manager failure after maximum retries.
domainOtherSideEportIsolation	The peer E port is isolated.
domainOverlapIsolation	There is a overlap in domains while attempting to connect two existing fabrics.
elpFailureClassFParamErr	Isolated for ELP failure due to class F parameter error.
elpFailureClassNParamErr	Isolated for ELP failure due to class N parameter error.
elpFailureInvalidFlowCTLParam	Isolated for ELP failure due to invalid flow control parameter.
elpFailureInvalidPayloadSize	Isolated for ELP failure due to invalid payload size.
elpFailureInvalidPortName	Isolated for ELP failure due to invalid port name.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 35-2 Device Manager Tooltip Definitions (continued)**

Tooltip	Definition
elpFailureInvalidSwitchName	Isolated for ELP failure due to invalid switch name.
elpFailureInvalidTxBBCredit	Isolated for ELP failure due to invalid transmit B2B credit.
elpFailureIsolation	During a port initialization the prospective Interconnect_Ports find incompatible link parameters.
elpFailureLoopbackDetected	Isolated for ELP failure due to loopback detected.
elpFailureRatovEdtovMismatch	Isolated for ELP failure due to R_A_TOV or E_D_TOV mismatch.
elpFailureRevMismatch	Isolated for ELP failure due to revision mismatch.
elpFailureUnknownFlowCTLCode	Isolated for ELP failure due to invalid flow control code.
ePortProhibited	Port down because FICON prohibit mask in place for E/TE port.
eppFailure	Trunk negotiation protocol failure after maximum retries.
errorDisabled	The port is not operational due to some error conditions that require administrative attention.
escFailureIsolation	During a port initialization the prospective Interconnect_Ports are unable to proceed with initialization as a result of Exchange Switch Capabilities (ESC).
fabricBindingDBMismatch	fabric binding active database mismatch with peer.
fabricBindingDomainInvalid	Peer domain ID is invalid in fabric binding active database.
fabricBindingNoRspFromPeer	Fabric binding no response from peer.
fabricBindingSWWNNotFound	Peer switch WWN not found in fabric binding active database.
fcipPortAdminCfgChange	FCIP port went down due to configuration change.
fcipPortKeepAliveTimerExpire	FCIP port went down due to TCP keep alive timer expired.
fcipPortMaxRetx	FCIP port went down due to max TCP retransmissions reached the configured limit.
fcipPortPersistTimerExpire	FCIP port went down due to TCP persist timer expired.
fcipPortSrcAdminDown	FCIP port went down because the source ethernet link was administratively shutdown.
fcipPortSrcLinkDown	FCIP port went down due to ethernet link down.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 35-2 Device Manager Tooltip Definitions (continued)**

Tooltip	Definition
fcipSrcModuleNotOnline	FCIP port went down due to source module not online.
fcipSrcPortRemoved	FCIP port went down due to source port removal.
fcotChksumErr	FSP SPROM checksum error.
fcotNotPresent	SFP (GBIC) not present.
fcotVendorNotSupported	FSP (GBIC) vendor is not supported.
fcspAuthenfailure	Fibre Channel security protocol authorization failed.
ficonBeingEnabled	FICON is being enabled.
ficonNoPortnumber	No FICON port number.
ficonNotEnabled	FICON not enabled.
ficonVsanDown	FICON VSAN is down.
firstPortNotUp	In a over subscribed line card, first port cannot be brought up in E mode when the other ports in the group are up.
firstPortUpAsEport	In a over subscribed line card, when the first port in a group is up in E mode, other ports in that group cannot be brought up.
hwFailure	Hardware failure.
incomAdminRxBBCreditPerBuf	Disabled due to incompatible admin port rxbbcredit, performance buffers.
incompatibleAdminMode	Port admin mode is incompatible with port capabilities.
incompatibleAdminRxBBCredit	Receive BB credit is incompatible.
incompatibleAdminRxBufferSize	Receive buffer size is incompatible.
incompatibleadminSpeed	Port speed is incompatible with port capabilities.
initializing	The port is being initialized.
interfaceRemoved	Interface is being removed.
invalidAttachment	Invalid attachment.
invalidConfig	This port has a misconfiguration with respect to port channels.
invalidFabricBindExh	Invalid fabric binding exchange.
linkFailCreditLoss	Link failure due to excessive credit loss indications.
linkFailCreditLossB2B	Link failure when link reset (LR) operation fails due to queue not empty.
linkFailDebounceTimeout	Link failure due to re-negotiation failed.
linkFailLineCardPortShutdown	Link failure due to port shutdown.
linkFailLinkReset	Link failure due to link reset.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 35-2 Device Manager Tooltip Definitions (continued)**

<b>Tooltip</b>	<b>Definition</b>
linkFailLIPF8Rcvd	Link failure due to F8 LIP received.
linkFailLIPRcvdB2B	Link failure when loop initialization (LIP) operation fails due to non empty receive queue.
linkFailLossOfSignal	Link failure due to loss of signal.
linkFailLossOfSync	Link failure due to loss of sync.
linkFailLRRcvdB2B	Link failure when link reset (LR) operation fails due to non-empty receive queue.
linkFailNOSRcvd	Link failure due to non-operational sequences received.
linkFailOLSRcvd	Link failure due to offline sequences received.
linkFailOPNyRETB2B	Link failure due to open primitive signal returned while receive queue not empty.
linkFailOPNyTMOB2B	Link failure due to open primitive signal timeout while receive queue not empty.
linkFailPortInitFail	Link failure due to port initialization failure.
linkFailPortUnusable	Link failure due to port unusable.
linkFailRxQOverflow	Link failure due to receive queue overflow.
linkFailTooManyINTR	Link failure due to excessive port interrupts.
linkFailure	Physical link failure.
loopbackDiagFailure	Loopback diagnostics failure.
loopbackIsolation	Port is connected to another port in the same switch.
noCommonVsanIsolation	Trunk is isolated because there are no common vsans with peer.
none	No failure.
nonParticipating	During loop initialization, the port is not allowed to participate in loop operations
offline	Physical link is in offline state as defined in the FC-FS standards.
ohmsExtLBTst	Link suspended due to external loopback diagnostics failure.
other	Undefined reason.
parentDown	The physical port to which this interface is bound is down.
peerFCIPPortClosedConnection	Port went down because peer FCIP port closed TCP connection.
peerFCIPPortResetConnection	Port went down because the TCP connection was reset by the peer FCIP port.
portBindFailure	Port got isolated due to port bind failure.
portBlocked	Port blocked due to FICON.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

**Table 35-2 Device Manager Tooltip Definitions (continued)**

Tooltip	Definition
portChannelMembersDown	No operational members.
portFabricBindFailure	Port isolated due to fabric bind failure.
portGracefulShutdown	Port shutdown gracefully.
portVsanMismatchIsolation	An attempt is made to connect two switches using non-trunking ports having different port VSANs.
rcfInProgress	An isolated xE_port is transmitting a reconfigure fabric, requesting a disruptive reconfiguration in an attempt to build a single, non-isolated fabric. Only the Interconnect_Ports can become isolated.
srcPortNotBound	No source port is specified for this interface.
suspendedByMode	Port that belongs to a port channel is suspended due to incompatible operational mode.
suspendedBySpeed	Port that belongs to a port channel is suspended due to incompatible operational speed.
suspendedByWWN	Port that belongs to a port channel is suspended due to incompatible remote switch WWN.
swFailure	Software failure.
tooManyInvalidFLOGIs	Suspended due to too many invalid FLOGIs.
tovMismatch	Link isolation due to TOV mismatch
trunkNotFullyActive	Some of the VSANs which are common with the peer are not up.
upgradeInProgress	Line card upgrade in progress.
vsanInactive	Port VSAN is inactive. The port becomes operational again when the port VSAN is active.
vsanMismatchIsolation	This VSAN is not configured on both sides of a trunk port.
zoneMergeFailureIsolation	The two Interconnect_Ports cannot merge zoning configuration after having exchanged merging request for zoning.
zoneRemoteNoRespIsolation	Isolation due to remote zone server not responding.

## Can I upgrade without losing my map settings?

When you upgrade from one version of Fabric Manager to another, there is a way to prevent the loss of map settings (enclosure names, placement on the map, etc.)

The MDS 9000/db directory contains subfolders for each user (and one for fmserver). In these subfolders are files for all discovered fabrics (\*.dat) and maps (\*.map). These are upgradable between versions. If you need to clear the fabric cache, you should first export the enclosures to a file to avoid losing them. Everything else aside from enclosures and map coordinates are stored on the switch. The preferences, last opened, and site\_ouis.txt format doesn't change from release to release.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## How can I preserve historical data when moving Fabric Manager server to a new host?

To preserve your data when moving Fabric Manager Server to a new host, follow these steps:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Copy the <code>cisco_mds9500/pm</code> directory from the old host to the new host. Place it in the MDS 9000 directory (on a Windows PC, the default installation location for this directory is <code>C:\Program Files\Cisco Systems\MDS 9000</code> ). |
| <b>Step 2</b> | On the new host, run <b>PMUpgrade.bat</b> from the MDS 900\bin folder. This creates some new files and a new directory structure. There is a directory for each switch for which you have collected data.  |
| <b>Step 3</b> | To continue data collection on a specific switch, copy the <b>db</b> subfolder from that switch's folder to the <b>pm</b> folder.  |
| <b>Step 4</b> | On the new host, restart the Performance Manager Service (Windows) or Daemon (Unix). You can use the <b>bin/PM.bat</b> file to do this, or you can choose <b>Performance &gt; Collector &gt; Restart</b> from the Fabric manager menu.                   |
| <b>Step 5</b> | Export the enclosures to a file.   |
| <b>Step 6</b> | Reimport the enclosers on the new host.  |
| <b>Step 7</b> | Be sure to turn off the original service on the old host.  |
- 

## Are there any restrictions when using fabric manager across FCIP?

Fabric Manager will work with no restriction across an FCIP tunnel, as long as the tunnel is up. However, Fabric Manager cannot automatically discover a Cisco SN5428 mgmt IP address in the fabric. For that switch, it will display a red slash through an FCIP device because of a timeout error. It will still see all targets, initiators, and ISLs attached to a Cisco SN5428 (or any other switch) as long as they appear in the name server or FSPF.

To work around this, you can manually enter the IP address in the Switches table, and click Apply. If the community string is correct, the red slash will go away. Even if the community string is incorrect, double-clicking on the Cisco SN5428 will launch the web tool.

## I see "Please insure that FM server is running on localhost."

You may see this error message if you cannot connect to the fabric and your PC has multiple network interface cards. The problem may be that Fabric Manager is trying to communicate through the wrong interface (you can verify this by checking the FMServer.log file).

Generally it is best to let Fabric Manager choose the interface on startup. If you are getting the above error, something may have gone wrong.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

To reset Fabric Manager so that it chooses the interface next time it starts, do the following:

- 
- |               |  |
|---------------|--|
| <b>Step 1</b> | Open the server.properties file in the Fabric Manager installation directory. On a Windows platform, this file is in C:\Program Files\Cisco Systems\MDS 9000 by default. |
| <b>Step 2</b> | Comment out the <b>snmp.localaddress</b> line.   |
| <b>Step 3</b> | Save and exit the file.  |
| <b>Step 4</b> | Restart Fabric Manager.  |
- 

**Note**

There are some cases where you would not want to do this, and should manually select the interface that Fabric Manager uses. For more information, see the [“How can I run Cisco Fabric Manager if I have multiple interfaces?”](#) section on page 35-21.

---

## How can I run Cisco Fabric Manager if I have multiple interfaces?

If your PC has multiple interfaces (NICs), the four Cisco Fabric Manager applications detect these interfaces automatically (ignoring loopback interfaces). Fabric Manager Client and Device Manager detect all interfaces on your PC each time you launch them, and allow you to select one. Fabric Manager Server and Performance Manager detect on initial install, and allows you to select one. You are not prompted again to choose an interface with these two applications.

There may be circumstances where you will want to change the interface you are using. For example:

- If you add an interface after you have installed Fabric Manager Server and/or Performance Manager
- If you decide to use a different interface than the one you initially selected
- If for any reason one of the Cisco Fabric Manager applications did not detect multiple interfaces

Refer to the following sections, depending on which application you want to recognize the interface.

- [Manually specifying an interface for Fabric Manager Server, page 35-21](#)
- [Manually specifying an interface for Fabric Manager Client or Device Manager, page 35-22](#)

## Manually specifying an interface for Fabric Manager Server

To specify an interface for Fabric Manager Server (including Performance Manager and Fabric Manager Web Services), follow these steps:

- 
- |               |   |
|---------------|---|
| <b>Step 1</b> | Go to the MDS 9000 folder. On a Windows platform, this folder is at C:\Program Files\Cisco Systems\MDS 9000 by default. |
| <b>Step 2</b> | Edit the server.properties file with a text editor.   |
| <b>Step 3</b> | Scroll until you find the line snmp.localaddress  |
| <b>Step 4</b> | If the line is commented, remove the comment character.   |

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Step 5** Set this value to the IP address or interface name of the NIC you want to use.
  - Step 6** Save the file.
  - Step 7** Stop and restart Fabric Manager Server.
- 

## Manually specifying an interface for Fabric Manager Client or Device Manager

To specify an interface for the Fabric Manager Client or Device Manager, follow these steps:

- Step 1** Go to the MDS 9000/bin folder. On a Windows platform, this folder is at C:\Program Files\Cisco Systems\MDS 9000 by default.
  - Step 2** Edit the DeviceManager.bat file or the FabricManager.bat file.
  - Step 3** Scroll to the line that begins with set JVMARGS=
  - Step 4** Add the parameter -Device Managerds.nmsaddress=ADDRESS, where ADDRESS is the IP address or interface name of the NIC you want to use.
  - Step 5** Save the file and relaunch Fabric Manager Client or Device Manager.
- 

## How can I configure an HTTP proxy server?

If your network uses a proxy server for HTTP requests, make sure the Java Web Start Application Manager is properly configured with the IP address of your proxy server.

To configure a proxy server in the Java Web Start Application Manager, follow these steps:

- Step 1** Launch the Java Web Start application.
  - Step 2** Select **File > Preferences** from the Java WebStart Application Manager.
  - Step 3** Click the **Manual** radio button and enter the IP address of the proxy server in the **HTTP Proxy** field.
  - Step 4** Enter the HTTP port number used by your proxy service in the **HTTP Port** field.
  - Step 5** Click **OK**.
-



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## How can I clear the topology map?

If you have a switch that you have removed from the fabric, there will be a red X through the switch's icon. You can clear this information from the Fabric Manager client, or from the Fabric Manager server (which will clear the information for all clients) without having to reboot the switch.

To clear information from topology maps, follow these steps:

---

**Step 1** In the Fabric pane, click on the **Refresh Map** icon.

This clears the information from the client.

**Step 2** From the Server menu, click **Purge**.

This clears the information from the server.



**Caution**

Any devices not currently accessible (may be offline) will be purged.

---

## Can I use Fabric Manager in a mixed software environment?

You can use Fabric Manager version 2.0(x) to manage a mixed fabric of Cisco MDS 9000 switches. Certain 2.0 feature tabs will be empty for any switches running a software version that does not support those features.

## I get an error when launching Fabric Manager

If you get the following error:

```
An error occurred while launching the application Fabric Manager.

download error:corrupted jar file at <ipaddress>\Device Managerboot.jar
```

(Where <ipaddress> is that of the switch)

The error message you are getting indicates that the Java Web Start cache is corrupted. You can try clearing your Java Web Start cache first. To clear the Cache either run Java Web Start (from the Programs menu) and under the **preferences** select **clear cache**. Or do it manually by first making sure all Fabric Manager or Device Manager instances are closed and then deleting .javaws/cache. In the newer JREs this directory is created under Documents and Settings\USERNAME and in the older ones it used to be under Program Files\Java Web Start.

You can also browse beneath the cache folder and delete the offending IPAddress folder (e.g. cache/http/D10.0.0.1).

Also, check to make sure that the host is not running a virus checker / java blocker?

Also you can run the un-install program, then deleting .cisco\_mds directory. Then re-install Fabric Manager.

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## Can I search for devices in a fabric?

In Fabric Manager, it is possible to search for one or more devices by different attributes, including pWWN.

To perform a search, follow these steps:

- 
- Step 1** Right click on the map and select the **Find Elements** menu item from the popup menu.  
You see the Find Elements dialog box.
  - Step 2** Select **End Device** in the first combo-box.
  - Step 3** Select **Port WWN** in the second combo-box.  
You can enter part of the WWN and use a wildcard (\*) character (e.g., `***fb*f8`).
  - Step 4** Click on **Find in map**.  
You see the devices highlighted in the Fabric pane. You can right click on any device to see the attributes for that devices. You can also select the links leading to those devices to see the attributes for those links.
- 

## Do I need a license of Fabric Manager Server for each switch in the fabric?

No.

You must install a Cisco MDS 9000 Family Cisco Fabric Manager Server package on at least one switch in each fabric where you intend to manage switches, if you intend to use the enhanced management capabilities the license package provides. You must also license all switches you plan to monitor with the Performance Manager (historical performance monitoring) feature. Failure to license all switches can prevent effective use of the Flow performance monitoring, so it is recommended to license all switches in each fabric managed by Cisco Fabric Manager Server.

You are free to try Cisco Fabric Manager Server capabilities prior to installing a license, but the those extended functions will stop working after the 120-day grace period expires. Standard Cisco Fabric Manager configuration and management capabilities will continue to be accessible without any licensed switches after the grace period expires.

## How can I manage multiple fabrics?

To monitor and manage multiple fabrics, you must persist one or more fabrics. Do this by checking the **Persist** checkbox on the **Server>Admin** dialog Fabric tab. You must also use switches running SAN-OS Release 1.3.x or greater in both fabrics, and you must use the same user/password on both fabrics. Both fabrics must not be physically connected.

## License expiration causes orange X through switch

If you are using a licensed feature and that license is allowed to expire, Fabric Manager shows a license violation, and an orange X is placed through the switch on the Fabric Manager map.

To clear the license violation message and the orange X, stop the Cisco Fabric Manager service on the host, and restart it again.



## GUI/CLI Usage Chart

In the table below, an X indicates the possible ways the procedure can be performed (using the CLI, using Fabric Manager, or using Device Manager).

### Procedures

Procedure Category	Procedure	CLI	FM	DM
General OS	Basic Switch Configuration	X		
	Terminal Settings	X		
	File System Commands	X		X
	Displaying File Contents	X		X
Licenses	Installing Licenses	X	X	X
	Uninstalling Licenses	X	X	X
	Updating Licenses	X		
	Moving licences between switches	X		
Initial Configuration	Starting a Switch (Initial Setup)	X		
	Accessing the Switch	X	X	
	Assigning a Switch Name	X	X	X
	NTP Configuration	X	X	X
	Configuring the Management Interface	X	X	X
	Configuring the Default Gateway	X	X	X
	Configuration File Manipulation	X	X	X
	Downgrading from a Higher Release	X	X	
	Accessing Remote File Systems	X	X	X
	Configuring Console Settings	X		
	Configuring COM1 and Modem Settings	X		
	CDP Configuration	X	X	X

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Procedure Category	Procedure	CLI	FM	DM
High Availability	Initiating a Switchover	X	X	X
	Synchronizing Supervisor Modules	X	X	X
	Copying Images to the Standby Supervisor	X		X
Software Images	Performing an Automated Upgrade	X	X	
	Performing a Manual Upgrade	X		
	Recovering a Corrupted Bootflash	X		
Hardware	Verifying the Status of a Module	X	X	X
	Checking the State of a Module	X	X	X
	Connecting to a Module	X	X	X
	Reloading Modules	X	X	X
	Preserving Module Configuration	X		
	Purging Module Configuration	X		
	Powering Off Switching Modules	X	X	X
	Upgrading EPLD Images	X		
	Displaying EPLD Versions	X	X	X
	Displaying Switch Hardware Inventory	X	X	X
	Displaying the Switch Serial Number	X	X	X
	Displaying Environment Information	X	X	X
CFS	Enabling CFS for an Application	X	X	X
	Enabling or disabling CFS per switch	X	X	X
	Locking the Fabric	X	X	X
	Committing Changes	X	X	X
	Discarding Changes	X	X	X
	Saving the Configuration	X	X	X
	Clearing a Locked Session	X	X	X
VSAN	Static VSAN Configuration	X	X	X
	Dynamic VSAN Configuration (DPVM)	X	X	
Interfaces	Fibre Channel Interfaces	X	X	X
	Management Interface Configuration	X	X	X
	Displaying the ALPA Cache Contents	X		
	Clearing the ALPA Cache	X		
	VSAN Interface Configuration	X	X	X
	Common Information Model (CIM) Configuration	X		
Trunking	Enabling or Disabling the Trunking Protocol	X	X	X
	Configuring the Trunk Mode	X	X	X
	Configuring An Allowed List of VSANs	X	X	X

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Procedure Category	Procedure	CLI	FM	DM
PortChannels	PortChannel Creation and Configuration	X	X	X
	Enabling and Configuring Autocreation	X	X	X
Zones	Configuring a Zone	X	X	X
	Alias Configuration	X	X	X
	Zone Set Creation	X	X	X
	Configuring the Default Zone Policy	X	X	X
	Importing and Distributing Zone Sets	X	X	X
	Configuring a LUN-Based Zone	X	X	X
	Configuring Read-Only Zones	X	X	X
	Enabling Enhanced Zoning	X	X	X
	Configuring DDAS	X		
Inter-VSAN Routing	IVR Topology Configuration	X	X	
FLOGI, Name Server, FDMI, and RSCN	Displaying FLOGI Details	X	X	X
	Registering Name Server Proxies	X	X	X
	Displaying FDMI	X	X	X
	RSCN Statistics	X	X	X
Security	Configuring RADIUS	X	X	X
	Configuring TACACS+	X	X	X
	Configuring Server Groups	X	X	X
	Configuring AAA Services	X	X	X
	Role-Based Authorization	X	X	X
	Configuring User Accounts	X	X	X
	SNMP Security	X	X	X
	Configuring Accounting Services	X	X	X
	Configuring SSH Services	X	X	X
	Recovering Administrator Password	X		
	DHCHAP Configuration	X	X	X
	Port Security Configuration	X	X	X
	AutoLearning	X	X	X
Intelligent Storage Services	Configuring Fibre Channel Write Acceleration	X	X	
	Configuring SCSI Flow Statistics	X	X	
	Configuring NASB	X	X	
	Configuring SANTap	X	X	

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Procedure Category	Procedure	CLI	FM	DM
Fibre Channel Routing	FSPF Global Configuration	X	X	X
	FSPF Interface Configuration	X	X	X
	Configuring Fibre Channel Routes	X	X	X
	Broadcast Routing	X	X	X
	Multicast Routing	X	X	X
	In-Order Delivery	X	X	X
	Flow Statistics Configuration	X	X	X
IP Services	Traffic Management Services	X		X
	Management Interface Configuration	X	X	X
	Default Gateway Configuration	X	X	X
	Default Network Configuration	X	X	X
	IP Access Control Lists	X	X	X
	IPFC Configuration	X	X	X
	Configuring IP Static Routes	X	X	X
	Overlay VSAN Configuration	X		X
	Multiple VSAN Configuration	X		X
	VRRP Configuration	X	X	X
	DNS Server Configuration	X	X	X
FICON	FICON Configuration	X	X	X
	Running Configuration Automatic Save	X		X
	Binding Port Numbers to PortChannels	X	X	X
	Binding Port Numbers to FCIP Interfaces	X	X	X
	Configuring FICON Ports	X		X
	FICON Configuration File Manipulation	X		X
	Port Swapping	X		X
	Clearing FICON Device Allegiance	X		X
	CUP In-Band Management	X		X
	Fabric Binding Configuration	X	X	X
IPsec Network Security	IKE Configuration	X	X	X
	IPsec Configuration	X	X	X
IP Storage	Configuring Gigabit Ethernet Interfaces	X	X	X
	Configuring FCIP	X	X	X
	Configuring iSCSI	X	X	X
	Configuring Storage Name Services (iSNS)	X	X	X
Call Home	Call Home Configuration	X	X	X

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Procedure Category	Procedure	CLI	FM	DM
Domain Parameters	Domain Configuration	X	X	X
	Switch Priority	X	X	X
	Allowed Domain ID Lists	X	X	X
	Persistent FC IDs	X	X	X
Traffic Management	FCC Configuration	X	X	X
	QoS	X	X	X
	Control Traffic	X	X	X
	Data Traffic	X	X	X
Port Tracking	Enabling Port Tracking	X	X	X
	Configuring Linked Ports	X	X	X
FCIP Traffic Performance	Tuning Configuration	X		
Scheduling Tasks	Scheduler Configuration	X		
System Messages	System Message Logging Configuration	X	X	X
Discovering SCSI Targets	SCSI LUN Discovery	X	X	X
SPAN	Configuring SPAN	X		X
	Monitoring Traffic	X		X
	Remote SPAN	X		
Advanced Configuration	fcTimer	X	X	X
	fctrace	X	X	
	fcping	X	X	
	Configuring World Wide Names	X	X	X
	Configuring a Secondary MAC Address	X	X	X
Fabric Configuration Servers	FCS Configuration	X		X
System Processes and Logs	Displaying System Processes	X		X
	Displaying System Status	X		X
	Clearing the Core Directory	X		
	Displaying Core Status	X		X
	System Health Initiation	X		
	Loopback Test Configuration Frequency	X		
	Hardware Failure Action	X		
	Tests for a Specified Module	X		
	Clearing Previous Error Reports	X		
	Performing Internal Loopbacks	X		X
	Performing External Loopbacks	X		X

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***





## Interface Nonoperational Reason Codes

If the administrative state for an interface is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in [Table B-1](#).

**Table B-1** Reason Codes for Nonoperational States

Reason Code	Description	Applicable Modes
Link failure or not connected	Physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	Cisco MDS SAN-OS waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state.  To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: <ul style="list-style-type: none"> <li>Configuration failure.</li> <li>Incompatible buffer-to-buffer credit configuration.</li> </ul> To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shutdown or enable the interface.	

**Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).**

**Table B-1 Reason Codes for Nonoperational States (continued)**

Reason Code	Description	Applicable Modes
Isolation due to ELP failure	Port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	Port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to other side E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	



---

## Numerics

16-port modules

BB\_credits [18-8](#)

32-port modules

BB\_credits [18-8](#)

configuration guidelines [18-5](#)

---

## A

AAA

description [26-1](#)

distributing with CFS (procedure) [26-11](#)

enabling distribution with CFS (procedure) [26-10](#)

local authentication [26-11](#)

RADIUS [26-5](#)

TACACS+ [26-7](#)

access control

iSCSI enforcement [20-17](#)

iSCSI mechanisms [20-16](#)

Access Control Lists. See ACLs

accounting [26-3](#)

ACLs

associating with an interface (procedure) [27-6](#)

configuration guidelines [27-1](#)

creating complex ACLs (procedure) [27-5](#)

creating with IP-ACL Wizard (procedure) [27-4](#)

deleting (procedure) [27-7](#)

removing from an interface (procedure) [27-6](#)

administrator passwords, recovering [24-7](#)

advanced mode in Fabric Manager client [3-2](#)

AFIDs

configuring (procedure) [16-9](#)

default (procedure) [16-9](#)

description [16-4](#)

aliases

creating zones with (procedure) [15-6](#)

deleting (procedure) [15-13](#)

description [15-6](#)

viewing (procedure) [15-7](#)

authentication, authorization, and accounting. See AAA

autonomous fabric ID

See AFIDs

AutoNotify [33-1](#)

auto port mode

description [18-4](#)

interface configuration [18-1](#)

auto-topology

configuration guidelines [16-5](#)

IVR [16-4](#)

modifying (procedure) [16-9](#)

---

## B

BB\_credits

16-port modules [18-8](#)

32-port modules [18-8](#)

configuring [18-8](#)

reason codes [B-1](#)

Berkeley Packet Filter. See BPF

bootflash

file system [10-1](#)

space requirements [10-2](#)

BPF

library [34-11](#)

See also libpcap freeware

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## B ports

- enabling (procedure) [19-17](#)
- functionality [18-4](#)
- interface modes [18-1, 18-4](#)
- interoperability mode [19-15](#)
- SAN extenders [19-16](#)

bridge port. See B port

buffer-to-buffer credits. See BB\_credits

## C

### Call Home

- alerts [33-3](#)
- alerts (procedure) [33-3](#)
- configuring (procedure) [33-2](#)
- description [33-1](#)
- device ID format [33-5](#)
- event triggers [33-3](#)
- message contents [33-5](#)
- message severity levels [33-3](#)
- profiles [33-2](#)
- profiles (procudure) [33-3](#)

### CFS

- committing changes (procedure) [12-5](#)
- description [12-1](#)
- disabling (isolating) or enabling distribution on a switch (procedure) [12-6](#)
- enabling (procedure) [12-3](#)
- example using Device Manager [12-9](#)
- example using Fabric Manager [12-7](#)
- MDS SAN-OS features [12-1](#)
- merge support (procedure) [12-8](#)
- merging [12-7](#)

Cisco.com IDs [33-2](#)

Cisco Fabric Services. See CFS

### Cisco Traffic Analyzer

- configuring with Performance Manager (procedure) [8-5](#)
- description [8-3](#)

installing (procedure) [8-4](#)

claim certificate [9-2](#)

### common roles

- creating (procedure) [24-2](#)
- deleting (procedure) [24-3](#)
- description [24-2](#)
- editing rules (procedure) [24-3](#)

### CompactFlash

- devices [10-8](#)
- disk [10-1](#)

### configuraiton files

- saving (procedures) [11-1](#)

### configuration files

- copying (procedure) [11-2](#)

Control Unit Port. See CUP

counted licenses [9-2](#)

cross-VSAN communication [16-13](#)

### crypto map

- any keyword [28-12](#)
- applying to interface [28-16](#)
- AutoPeer option [28-14](#)
- configuration guidelines [28-9](#)
- entries [28-13](#)
- managing entries (procedure) [28-16](#)
- mirror image [28-10](#)

### CUP

- description [21-1](#)
- management [21-20](#)

## D

### databases

See zone databases

databases. See zone databases

### default zones

- description [15-14](#)
- interoperability [23-5](#)
- setting policy (procedure) [15-15](#)
- See also zones

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- destination IDs
    - fctrace [23-2](#)
    - PortChannels [17-1](#)
  - detachable tables [3-12](#)
  - device alias [1-7, 2-8](#)
  - device grouping (procedure) [3-14](#)
  - device IDs
    - Call Home format [33-6](#)
  - Device Manager
    - description [4-1](#)
    - launching (procedure) [4-2](#)
    - port status [4-6](#)
    - preferences [4-8](#)
    - tabs [4-5](#)
    - using interface (figure) [4-3](#)
    - viewing license information [9-10](#)
  - DHCHAP
    - authentication modes [29-4](#)
    - changing group list (procedure) [29-6](#)
    - changing hash algorithm (procedure) [29-5](#)
    - changing timeout value (procedure) [29-8](#)
    - configuring [29-3](#)
    - configuring authentication mode (procedure) [29-5](#)
    - configuring password [29-6](#)
    - configuring password (procedure) [29-7](#)
    - configuring remote password (procedure) [29-7](#)
    - description [29-2](#)
    - enabling (procedure) [29-3](#)
    - using with Cisco MDS SAN-OS [29-3](#)
  - digital signature algorithm. See DSA key pairs
  - discovering a fabric, best practices [7-3](#)
  - disruptive upgrades [10-3](#)
  - documentation
    - obtaining [xxxv](#)
    - related [xxxiv](#)
  - domain IDs
    - failure [B-2](#)
    - interoperability [23-5](#)
    - non-unique and IVR NAT [16-3](#)
    - unique [16-6](#)
  - domain overlap [B-2](#)
  - DPVM
    - autolearn entries [14-3](#)
    - databases [14-2](#)
    - description [14-1](#)
    - modifying databases (procedure) [14-5](#)
    - requirements [14-2](#)
    - using DPVM Setup Wizard (procedure) [14-4](#)
  - Dynamic Port VSAN Membership. See DPVM
- 
- ## E
- EISL, PortChannel links [17-1](#)
  - ELP failure [B-2](#)
  - enclosures [31-2](#)
  - enclosures (procedure) [3-14](#)
  - Enterprise package [9-3](#)
  - E ports
    - 32-port guidelines [18-5](#)
    - classes of service [18-2](#)
    - configuring [19-18](#)
    - interface modes [18-1](#)
    - isolation [B-2](#)
    - recovering from isolation [15-15](#)
  - error disabled code, interface [B-1](#)
  - ESC failure [B-2](#)
  - Ethereal freeware
    - analyzer [34-8](#)
    - information [34-7](#)
  - Ethernet PortChannels [19-21](#)
  - evaluation license [9-2](#)
  - events
    - description [5-3](#)
    - Device Manager [31-10](#)
    - Fabric Manager [31-10](#)
    - Fabric Manager Web Services [31-10](#)
  - exchange IDs, load balancing [17-1, 23-2](#)
  - exchange link parameter failure. See ELP failure

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

expiry alerts, licenses [9-11](#)

exporting

Performance Manager reports as CSV [32-9](#)

Performance Manager reports as XML [32-8](#)

zone databases [15-15](#)

## F

fabric, removing from monitoring [2-8](#)

Fabric Analyzer

capture filters [34-11](#)

configuring [34-9](#)

defining display filters [34-11](#)

description [34-7](#)

displaying captured frames [34-10](#)

permitted filters [34-12](#)

fabric binding

activating (procedure) [21-23](#)

clearing statistics (procedure) [21-25](#)

compare to port security [21-20](#)

configuring (procedure) [21-22](#)

copying to configuration file (procedure) [21-23](#)

creating config database (procedure) [21-24](#)

deactivating (procedure) [21-23](#)

deleting from config database (procedure) [21-24](#)

description [21-20](#)

enabling (procedure) [21-22](#)

enforcing [21-21](#)

viewing active database (procedure) [21-24](#)

viewing EFMD statistics (procedure) [21-25](#)

viewing violations (procedure) [21-24](#)

Fabric Manager

description [1-1](#)

detachable tables [3-12](#)

downgrade to release 1.3(x) [1-9](#)

downgrade to release 2.x [1-9](#)

installation [1-5](#)

installation (procedure) [1-7](#)

integrating with other tools [1-11](#)

launching (procedure) [1-9](#)

running behind a firewall [1-11](#)

uninstalling [1-12](#)

upgrade [1-8](#)

viewing license information [9-8](#)

Fabric Manager Client

advanced mode [3-2](#)

description [3-1](#)

filtering [3-11](#)

icons (table) [3-5](#)

Information pane icons (table) [3-7, 3-9, 4-4](#)

launching [3-2](#)

setting preferences [3-12](#)

troubleshooting tools [3-16](#)

using interface (figure) [3-3](#)

wizards [3-15](#)

Fabric Manager graphics (table) [3-5](#)

Fabric Manager Server

authentication [7-2](#)

configuring Performance Manager (procedure) [2-4](#)

continuously monitoring a fabric (procedure) [2-7](#)

description [2-2](#)

installation overview [2-2](#)

installing Fabric Manager Web Services  
(procedure) [2-6](#)

licensing [2-3, 9-12](#)

properties file [2-8](#)

removing a fabric from monitoring (procedure) [2-8](#)

setting the seed switch (procedure) [2-4](#)

Fabric Manager Web Services

Admin tab [5-4](#)

authentication [7-4](#)

description [5-1](#)

events [5-3](#)

filter tree [5-2](#)

guest user [5-9](#)

installation (procedures) [5-5](#)

inventory [5-4](#)

launching (procedure) [5-7](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Performance tab [5-4](#)
- using with SSL [5-6](#)
- FC aliases, configuring zones [15-3](#)
- FC IDs
  - allocating [23-4](#)
  - allocating flat FC IDs [23-4](#)
  - configuring persistent FC IDs (procedure) [18-9](#)
  - configuring zones [15-2](#)
- FCIP
  - active mode [19-14](#)
  - checking trunk status (procedure) [19-10](#)
  - configuring peers [19-13](#)
  - configuring with FCIP Wizard (procedure) [19-5](#)
  - description [19-2](#)
  - Gigabit Ethernet ports [19-1](#)
  - high availability [19-18](#)
  - initiating an IP connection [19-14](#)
  - interfaces [19-3, 19-9](#)
  - link end points [19-3](#)
  - link failures [19-19](#)
  - links [19-3, 19-13](#)
  - modifying links and profiles (procedure) [19-9](#)
  - profiles [19-3, 19-8](#)
  - TCP connections [19-3](#)
  - verifying interfaces and Extended Link Protocol (procedure) [19-10](#)
- FCIP compression
  - configuring (procedure) [19-7](#)
  - description [19-5](#)
  - modifying (procedure) [19-11](#)
- FCIP tape acceleration
  - description [19-11](#)
  - enabling (procedure) [19-13](#)
- FCIP write acceleration
  - configuring (procedure) [19-7](#)
  - description [19-4](#)
  - modifying (procedure) [19-11](#)
- FCP, routing requests with iSCSI [20-4](#)
- fcping
  - description [23-2](#)
  - starting (procedure) [23-3](#)
- FC-SP
  - description [29-1](#)
  - enabling (procedure) [29-3](#)
  - enabling on ISLs [29-8](#)
- fctrace
  - description [23-2](#)
  - path discovery [23-2](#)
  - starting (procedure) [23-2](#)
- Fibre Channel over IP. See FCIP
- Fibre Channel PortChannels [19-21](#)
- Fibre Channel Security Protocol. See FC-SP
- Fibre Channel time out values [23-1](#)
- Fibre Channel write acceleration
  - configuration (procedure) [22-5](#)
  - description [22-4](#)
- FICON
  - allowing host control of switch (procedure) [21-13](#)
  - allowing host to configure FICON (procedure) [21-13](#)
  - calculating flow load balance (procedure) [21-26](#)
  - code-page option [21-12](#)
  - configuration files [21-16](#)
  - configure port blocking or port prohibiting (procedure) [21-15](#)
  - creating VSANs (procedure) [21-10](#)
  - description [21-1](#)
  - director history (procedure) [21-12](#)
  - displaying port configuration (procedure) [21-15](#)
  - displaying RLIR (procedure) [21-25](#)
  - enabling (procedure) [21-10](#)
  - fabric binding [21-20](#)
  - FC ID allocation [21-8](#)
  - FC ID last byte requirement [21-12](#)
  - managing configuration files (procedure) [21-18](#)
  - port blocking [21-14](#)
  - port numbering (table) [21-7](#)
  - port prohibiting [21-14](#)
  - port swap [21-18](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

port swap (procedure) [21-19](#)  
 uninstalled ports [21-8](#)  
 viewing port attributes (procedure) [21-16](#)  
 VSAN requirements [21-9](#)

File Transfer Protocol. See FTP

FL ports

classes of service [18-3](#)  
 fctrace [23-2](#)  
 interface modes [18-1](#)  
 nonparticipating code [B-2](#)

F ports

classes of service [18-2](#)  
 interface modes [18-1](#)

FSPF interoperability [23-5](#)

FSP not present [B-1](#)

FTP

Fx ports

32-port default [18-6](#)  
 interface modes [18-4](#)

## G

Gigabit Ethernet ports [19-1](#)  
 global alias [1-7, 2-8](#)  
 grace period [9-2](#)

## H

hard zoning

See also zones [15-14](#)

hard zoning, description [15-14](#)

HBA [20-3](#)

high availability

Ethernet PortChannel [20-23](#)  
 features [20-20](#)  
 iSCSI statically imported targets [20-10](#)  
 licensing [9-4](#)  
 PortChannels [17-1](#)

software upgrade [10-3](#)

VRRP [19-20, 20-22](#)

Host Bus Adaptor. See HBA

IDs

contract IDs [33-5](#)  
 image version and IDs [10-1](#)  
 serial IDs [33-6](#)  
 server IDs [33-6](#)  
 site IDs [33-5](#)

IKE

description [28-4](#)  
 supported transforms and algorithms [28-6](#)  
 terms [28-4](#)  
 viewing configuration [28-9](#)  
 Windows and Linux support [28-7](#)

system images

images. See kickstart images

inactive code, interface [B-1](#)

in-band management [1-5](#)

incremental licenses [9-2](#)

Information pane [3-7, 3-9, 4-4](#)

intelligent storage services

description [22-1](#)  
 disabling (procedure) [22-3](#)  
 disabling with force option [22-3](#)  
 enabling (procedure) [22-2](#)

interfaces

administrative states [18-5](#)  
 BB\_credits [18-8](#)  
 configuring BB\_credits or performance buffers (procedure) [18-8](#)  
 configuring Fibre Channel (procedure) [18-6](#)  
 configuring Gigabit Ethernet (procedure) [18-7, 19-2](#)  
 configuring trunking mode (procedure) [18-4](#)  
 description [18-1](#)  
 enabling or disabling (procedure) [18-7](#)



***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- managing attributes for ports [18-7](#)
- modes [18-1](#)
- nonoperational reason codes [B-1](#)
- operational state [18-5](#)
- performance buffers [18-8](#)
- reason codes [18-5](#)
- SFT types [18-9](#)
- states [18-5](#)
- internal bootflash
  - description [10-7](#)
  - Flash devices [10-7](#)
  - See also bootflash
- Internet storage name service. See iSNS
- interoperability [23-4](#)
- interop mode, configuring (procedure) [23-6](#)
- inter-VSAN routing. See IVR
- inventory
  - description [5-4](#)
  - Fabric Manager Client [31-3](#)
  - Fabric Manager Web Services [31-3](#)
- IP Access Control Lists. See ACLs
- IPFC interface configuration [18-9](#)
- IP filters
  - description [27-2](#)
  - using IP-ACL Wizard (procedure) [27-4](#)
- IPsec
  - compatibility [28-3](#)
  - configuring with FCIP Wizard (procedure) [19-6](#)
  - crypto map [28-9](#)
  - description [28-3](#)
  - enabling with FCIP Wizard (procedure) [28-7](#)
  - global lifetime values for SAs [28-17](#)
  - requirements [28-2](#)
  - SA establishment [28-14](#)
  - SA lifetime negotiation [28-15](#)
  - supported transforms [28-5](#)
  - terms [28-4](#)
  - transform sets [28-12](#)
  - viewing configuration (procedure) [28-9](#)
  - Windows and Linux support [28-7](#)
- IQN formats [20-8](#)
- iSCSI
  - as member of a zone (procedure) [20-16](#)
  - configuring trespass option (procedure) [20-11](#)
  - description [20-1 to 20-5](#)
  - discovery [20-17](#)
  - drivers [20-4](#)
  - enabling (procedure) [20-5](#)
  - Fibre Channel targets [20-2](#)
  - Gigabit Ethernet ports [19-1](#)
  - modifying VSANS assigned (procedure) [20-13](#)
  - overriding global interface authentication [20-18](#)
  - physical targets [20-2](#)
  - QoS (procedure) [20-19](#)
  - revert to primary port (procedure) [20-10](#)
  - routing requests [20-4](#)
  - session creation [20-17](#)
  - using iSCSI Wizard (procedure) [20-5 to 20-7](#)
  - virtual Fibre Channel hosts [20-11](#)
- iSCSI authentication
  - configuring (procedure) [20-18](#)
  - configuring RADIUS (procedure) [20-18](#)
  - description [20-17](#)
- iSCSI hosts
  - configuring accessibility [20-17](#)
  - multiple VSANs [20-13](#)
- iSCSI initiators
  - access list [20-17](#)
  - dynamic mapping [20-12](#)
  - identifying [20-12](#)
  - statically mapped (procedure) [20-13](#)
  - static mapping [20-12](#)
  - static mapping, manual [20-13](#)
  - static mapping, system assignment [20-13](#)
- iSCSI proxy initiators
  - configuring (procedure) [20-16](#)
  - description [20-14](#)
- iSCSI targets

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- access control [20-16](#)
- advertising [20-9](#)
- dynamic importing [20-7](#)
- dynamic importing (procedure) [20-8](#)
- secondary access [20-10](#)
- static importing [20-7](#)
- static importing (procedure) [20-9](#)
- ISL
  - PortChannel links [17-1](#)
  - statistics [32-2](#)
- iSNS
  - creating profile (procedure) [20-24](#)
  - description [20-23](#)
  - tagging a profile for an interface (procedure) [20-24](#)
- IVR
  - auto-topology [16-4](#)
  - border switch [16-3](#)
  - border switch, guidelines [16-7](#)
  - configuring (procedure) [16-10](#)
  - configuring zones and zone sets (procedure) [16-14](#)
  - default zone policy [16-13](#)
  - description [16-1](#)
  - domain ID guidelines [16-6](#)
  - edge switch [16-3](#)
  - edge VSAN [16-3](#)
  - Fibre Channel header modifications [16-3](#)
  - interoperability [16-7](#)
  - modifying [16-8](#)
  - path [16-2](#)
  - recovering the full topology (procedure) [16-17](#)
  - recovering the full zone database [16-16](#)
  - sharing resources [16-2](#)
  - terminology [16-2](#)
  - transit VSAN, guidelines [16-7](#)
  - using IVR Zone Wizard [16-7](#)
  - virtual domains [16-12 to 16-13](#)
  - VSAN topology [16-4](#)
  - zones (definition) [16-2](#)
  - zone set activation (procedure) [16-16](#)

- zone sets (definition) [16-2](#)
- IVR NAT
  - border switch, guidelines [16-5](#)
  - configuration guidelines [16-5](#)
  - description [16-3](#)
  - modifying (procedure) [16-9](#)
  - service groups, guidelines [16-6](#)
  - transit VSANs, guidelines [16-5](#)
  - using IVR Zone Wizard [16-7](#)
- IVR NAT auto-topology [16-4](#)
- IVR topology
  - activating (procedure) [16-12](#)
  - clearing manual entries [16-12](#)
  - creating manually [16-11](#)
  - enabling auto-topology [16-7](#)

---

## K

- kickstart images
  - description [10-1](#)
  - KICKSTART variable [10-1](#)

---

## L

- launching management software [1-9](#)
- libpcap freeware [34-7](#)
- license
  - enforcing [9-1](#)
  - Fabric Manager Server package [9-4](#)
  - factory-installed [9-5](#)
  - feature based [9-2](#)
  - high availability [9-4](#)
  - host ID [9-2](#)
  - installation options [9-5](#)
  - installing with License Wizard [9-7](#)
  - Mainframe package [9-4](#)
  - obtaining a license key file [9-6](#)
  - SAN extension package [9-4](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- Standard package [9-3](#)
- terminology [9-1](#)
- licenses
  - expiry alerts [9-11](#)
  - moving between switches [9-12](#)
  - uninstalling [9-10](#)
  - updating [9-11](#)
- link failures [B-1](#)
- load balancing
  - description [17-1](#)
  - FSPF [19-19](#)
  - PortChannel [19-19](#)
- local capture [34-9](#)
- log files, configuring [31-7](#)
- logging
  - severity levels [31-6](#)
  - system messages [31-4](#)
- Logical Domains pane [3-10](#)
- logical units. See LUs
- logs
  - Device Manager [31-9](#)
  - Fabric Manager Web Services [31-9](#)
  - RMON [33-16](#)
  - SNMP events [33-13](#)
- loop monitoring [23-4](#)
- loop port [23-4](#)
- LUN zoning
  - configuring [15-20](#)
  - description [15-19](#)
- LUs
  - exporting [20-11](#)
  - importing targets [20-7](#)

## M

- main menu [3-7](#)
- management protocols supported (table) [1-3](#)
- mapping
  - iSCSI hosts [20-11](#)

- PortChannels [19-21](#)
- module-based licensing [9-2](#)
- MPS-14/2 module
  - functions [28-1](#)
  - IPsec support [28-2](#)
- multiple fabrics [31-3](#)
- Multiprotocol Services module. See MPS-14/2 module

## N

- name server interoperability [23-6](#)
- NASB
  - configuring (procedure) [22-13](#)
  - description [22-11](#)
  - permitting default zoning (procedure) [22-12](#)
- native VSAN [16-2](#)
- network administrator [26-2](#)
- Network-Assisted Serverless Backup. See NASB
- network operator [26-2](#)
- new and changed information (table) [xxv](#)
- NL ports
  - fctrace [23-2](#)
  - interface modes [18-4](#)
  - zone enforcement [15-13](#)
- node-locked license [9-2](#)
- nondisruptive upgrades [10-3](#)
- N ports
  - fctrace [23-2](#)
  - zone enforcement [15-13](#)
- Nx ports, hard zoning [15-14](#)

## O

- offline code, interface [B-1](#)
- out-of-band management [1-5](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

## P

PAA-2 [8-3](#)

passive mode, IP connection [19-14](#)

passwords

creating strong passwords [24-5](#)

recovering [24-7](#)

peers

configuring [19-13](#)

configuring (procedure) [19-14](#)

performance

historical monitoring [32-2](#)

ISL statistics (procedure) [32-2](#)

monitoring in Device Manager (procedure) [32-1](#)

per-port monitoring (procedure) [32-2](#)

real-time monitoring [32-1](#)

Performance Manager

architecture [6-1](#)

authentication [7-3](#)

collections [32-3](#)

configuring flows and collections (procedure) [32-5](#)

configuring with Traffic Analyzer [8-5, 32-10](#)

creating a flow [32-3](#)

creating a flow (procedure) [32-3](#)

creating collections (procedure) [32-5](#)

data collection [6-2](#)

data interpolation [6-2](#)

description [1-3](#)

events [32-7](#)

exporting as CSV [32-9](#)

exporting as XML [32-8](#)

reports [32-6](#)

starting collections [32-6](#)

thresholds [32-4](#)

top 10 reports [32-7](#)

top 10 reports (procedure) [32-8](#)

using thresholds [6-2](#)

verifying collections [2-6](#)

performance reports, description [5-4](#)

permanent license [9-2](#)

Physical Attributes pane [3-10](#)

Port Analyzer Adapter 2. See PAA-2

PortChannels

description [17-1](#)

down state [B-2](#)

interfaces using iSCSI [20-9](#)

interoperability [23-5](#)

IQN formats [20-8](#)

load balancing [19-19](#)

modifying (procedure) [17-5](#)

trunking comparison [17-1](#)

using the PortChannel Wizard (procedure) [17-2](#)

port group [18-5](#)

port modes

auto [18-4](#)

IPS [19-1](#)

port security

\*(wildcard) [30-6](#)

activating with auto-learn (procedure) [30-3](#)

activating without auto-learning (procedure) [30-10](#)

auto-learn [30-1](#)

configuring (procedure) [30-3](#)

configuring manually [30-6](#)

configuring manually (procedure) [30-7](#)

copying active to config database (procedure) [30-11](#)

databases [30-8](#)

deleting entries from database (procedure) [30-7](#)

description [30-1](#)

displaying settings (procedure) [30-4](#)

displaying statistics (procedure) [30-4](#)

displaying violations (procedure) [30-4](#)

enable or disable auto-learning (procedure) [30-5](#)

enabling (procedure) [30-3](#)

enforcing [30-2](#)

example [30-5](#)

reactivating the database (procedure) [30-11](#)

preferences

Device Manager [4-8](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

Fabric Manager Client [3-12](#)  
 preshared key [26-6](#)  
 product authorization key [9-2](#)  
 proof of purchase [9-2](#)  
 protocol analysis [34-7](#)

## R

R\_A\_TOV time [B-1](#)  
 RADIUS  
   AAA solutions [26-1](#)  
   adding a server (procedure) [26-5](#)  
   description [26-5](#)  
   secret key [26-1](#)  
   setting global preshared key (procedure) [26-6](#)  
   setting preshared key [26-6](#)  
   specifying SNMPv3 on AAA servers [26-7](#)  
   vendor-specific attributes protocol options [26-6](#)  
 reason codes [B-1](#)  
 reconfigure fabric [B-1](#)  
 recovering databases. See zone databases [15-18](#)  
 recovering passwords [24-7](#)  
 redundancy  
   Ethernet PortChannels [19-21](#)  
   Fibre Channel PortChannels [19-21](#)  
 remote capture [34-9](#)  
 Remote Capture Protocol. See RPCAP  
 reserved words in user accounts [24-5](#)  
 RMON  
   defining an event (procedure) [33-15](#)  
   description [33-13](#)  
   enabling alarms (procedure) [33-14](#)  
   setting alarms (procedure) [33-13](#)  
   viewing alarms (procedure) [33-15](#)  
   viewing log (procedure) [33-16](#)  
 roles  
   additional [26-2](#)  
   common [24-2](#)  
   deleting (procedure) [24-3](#)

network administrator [26-2](#)  
 network operator [26-2](#)  
 RPCAP  
   Ethereal communication [34-8](#)

## S

SAN discovery [31-1](#)  
 SAN operating system. See SAN-OS  
 SAN-OS [10-1](#)  
 SANTap  
   configuring (procedure) [22-10](#)  
   description [22-7 to 22-10](#)  
   proxy mode-1 [22-9](#)  
   proxy mode-2 [22-10](#)  
   transparent mode [22-8](#)  
 SCSI, routing requests using iSCSI [20-4](#)  
 SCSI flow services  
   configuring (procedure) [22-4](#)  
   description [22-3](#)  
   Fibre Channel write acceleration [22-4](#)  
 SCSI flow statistics  
   clearing (procedure) [22-6](#)  
   description [22-5](#)  
   enabling (procedure) [22-6](#)  
 SCSI routing [20-1](#)  
 SD ports [18-3](#)  
 security control  
   local [26-1](#)  
   remote [26-1](#)  
 service groups  
   configuring (procedure) [16-9](#)  
   IVR [16-4](#)  
 SFPs, showing types [18-9](#)  
 shutdown state [18-5](#)  
 SMARTnet [33-2](#)  
 SNMP  
   access control [25-2](#)  
   access groups [25-5](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- adding community strings (procedure) [25-3](#)
  - community strings [25-2](#)
  - configuring event destinations [33-12](#)
  - configuring event security (procedure) [33-12](#)
  - configuring notifications (traps or informs) (procedure) [25-6](#)
  - deleting community strings (procedure) [25-3](#)
  - description [25-1](#)
  - event destinations [33-12](#)
  - events [33-11](#)
  - filtering events (procedure) [33-11](#)
  - notifications [25-6](#)
  - server contact [33-1](#)
  - synchronized to CLI [25-2](#)
  - trap receivers [33-12](#)
  - users with multiple roles (procedure) [25-5](#)
  - versions [25-1](#)
  - viewing event log [33-13](#)
  - SNMPv3 security features [25-2](#)
  - software images
    - description [10-1](#)
    - upgrade prerequisites [10-2](#)
    - upgrade requirements [10-2](#)
    - upgrading [10-1](#)
    - variables [10-1](#)
  - Software Installation Wizard (procedure) [10-4](#)
  - software upgrades [10-3](#)
  - soft zoning [15-13](#)
  - source IDs
    - Call Home event format [33-6](#)
    - frame loop back [23-2](#)
    - load balancing [17-1](#)
  - SPAN, monitoring traffic [8-2](#)
  - SSH
    - enabling (procedure) [24-7](#)
    - generating server key pair (procedure) [24-7](#)
    - host key pair [24-6](#)
    - overwriting key pair [24-7](#)
  - SSM [22-1](#)
  - disabling intelligent storage services (procedure) [22-3](#)
  - enabling intelligent storage services (procedure) [22-2](#)
  - provisioning [22-1](#)
  - provisioning (procedure) [22-2](#)
  - status bar [3-10](#)
  - Storage management solutions architecture [1-4](#)
  - Storage Services Module. See SSM
  - ST ports [18-1, 18-4](#)
  - supervisor and switching modules in Device Manager [4-7](#)
  - syslog. See system messages
  - syslog server
    - configuring [31-8](#)
    - verifying [31-9](#)
  - system images
    - description [10-1](#)
    - SYSTEM variable [10-1](#)
  - system messages
    - configuring a syslog server [31-8](#)
    - configuring logging [31-7](#)
    - description [31-4](#)
    - enabling for a feature (procedure) [31-7](#)
    - facilities and severity levels [31-4](#)
    - severity levels for console session [31-7](#)
    - viewing from Device Manager [31-9](#)
    - viewing from Fabric Manager Web Services [31-9](#)
- 
- ## T
- TACACS+
    - adding a server (procedure) [26-8](#)
    - description [26-7](#)
    - supported servers [26-9](#)
  - TCP connections
    - FCIP profiles [19-8](#)
    - specifying [19-14](#)
  - TCP ports recognized in ACLs [27-3](#)
  - technical support [xxxvii](#)
  - Telnet [24-6](#)
  - TE ports

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- classes of service [18-3](#)
- fctrace [23-2](#)
- interface modes [18-1](#)
- interoperability [23-5](#)
- recovering from isolation [15-15](#)
- Threshold Manager [33-13](#)
- TL ports
  - classes of service [18-3](#)
  - interface modes [18-1](#)
- topology map
  - customizations [31-2](#)
  - description [31-1](#)
  - enclosures [31-2](#)
  - mapping multiple fabrics [31-3](#)
  - saving custom layout (procedure) [31-2](#)
- TOV
  - interoperability [23-5](#)
  - ranges [23-1](#)
- Traffic Analyzer. See Cisco Traffic Analyzer
- transit VSANs
  - configuration guidelines [16-5](#)
  - description [16-3, 16-11](#)
- troubleshooting
  - analyzing switch health [34-3](#)
  - analyzing zone merge [34-13](#)
  - Cisco Fabric Analyzer [34-7](#)
  - locating other switches [34-15](#)
  - loopback tests [34-4](#)
  - monitoring oversubscription [34-16](#)
  - show tech support [34-14](#)
  - system messages [31-4](#)
  - testing end-to-end connectivity [34-6](#)
  - tools [34-1](#)
  - using Fabric Configuration tool (procedure) [34-5](#)
  - using ping [34-12](#)
  - using traceroute (procedure) [34-13](#)
  - with protocol analyzer [34-3](#)
  - with Traffic Analyzer [34-2](#)
- trunking interoperability [23-5](#)

## U

- UDP ports in ACLs [27-3](#)
- uninstalling management software [1-12](#)
- uninstalling permanent licenses [9-10](#)
- updating licenses [9-11](#)
- upgrades. See disruptive upgrades
- upgrades. See nondisruptive upgrades
- upgrading management software [1-8](#)
- user ID for authentication [26-2](#)
- users
  - creating (procedure) [24-5](#)
  - deleting (procedure) [24-6](#)
  - reserved words [24-5](#)
  - viewing (procedure) [24-6](#)

## V

- VE ports [19-2](#)
- virtual E ports. See VE ports
- virtual Fibre Channel hosts [20-11](#)
- virtual iSCSI targets [20-2](#)
- virtual ISL [19-3](#)
- virtual N port
  - dynamic mapping [20-12](#)
- volatile file system [10-8](#)
- VRRP
  - high availability [19-20](#)
  - IQN formats [20-8](#)
  - switchover [19-20](#)
- VSA. See RADIUS
- VSAN IDs
  - configuring FICON [21-3](#)
  - multiplexing traffic [18-3](#)
  - range [13-1](#)
- VSANs
  - configuring (procedure) [13-2](#)
  - default [13-2](#)
  - deleting (procedure) [13-3](#)

***Send documentation comments to [mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com).***

- description [13-1](#)
- editing policy or scope (procedure) [24-4](#)
- FC IDs [13-1](#)
- features [13-1](#)
- interop mode [23-5](#)
- IPFC interface and fctrace [23-2](#)
- isolated [13-2](#)
- mismatch [B-2](#)
- multiple zones [15-8](#)
- policy [24-3](#)
- trunking [17-1](#)
- trunking port [18-3](#)

---

## W

world wide names. See WWNs

### WWNs

- address pool [20-12](#)
- configuring [23-3](#)
- suspended connection [B-2](#)

---

## Z

### zone databases

- copying (procedure) [15-18](#)
- importing [15-15](#)
- migrating a non-MDS database [15-18](#)

### zone members

- adding to zone [15-5](#)
- default [15-14](#)
- displaying [15-6](#)

### zones

- access control [15-8](#)
- active [15-8](#)
- adding to zone set (procedure) [15-11](#)
- adding zone members [15-5](#)
- analyzing merge [34-13](#)
- assigning members [15-2](#)

- backing up and restoring (procedure) [15-22](#)
- cloning [15-12](#)
- configuring [15-2](#)
- configuring (procedure) [15-5](#)
- configuring a LUN-based zones (procedure) [15-20](#)
- creating read-only zones (procedure) [15-22](#)
- creating with aliases (procedure) [15-6](#)
- default policy [15-2](#)
- deleting (procedure) [15-13](#)
- description [15-1](#)
- edit full zone database [15-4](#)
- enforcing [15-13](#)
- fabric pWWNs [15-2](#)
- FC aliases [15-3](#)
- features [15-1](#)
- LUN zoning [15-19](#)
- port IDs [15-2](#)
- pWWNs [15-2](#)
- QoS or broadcast attributes (procedure) [15-19](#)
- read-only [15-21](#)
- traffic priority [15-18](#)
- See also default zones
- See also hard zoning
- See also soft zoning

zone sets

- activating (procedure) [15-12](#)
- adding zones (procedure) [15-11](#)
- cloning [15-12](#)
- configuration guidelines [15-8](#)
- considerations [15-8](#)
- creating (procedure) [15-11](#)
- deactivating (procedure) [15-12](#)
- deleting (procedure) [15-13](#)
- description [15-8](#)
- distributing [15-17](#)
- exporting (procedure) [15-17](#)
- importing (procedure) [15-16](#)
- propagating [15-17](#)